

Private Information and Unemployment Insurance

Nathaniel Hendren*

July, 2015

Abstract

I provide empirical evidence that private information limits the ability to insure the risk of job loss in private markets, and analyze the implications of this micro-foundation for normative analysis of optimal unemployment insurance. Using information contained in subjective probability elicitation, I show individuals have significant information about their chances of losing their job conditional on a wide range of observable information insurers could potentially use to price the insurance. I derive lower bounds that suggest individuals would need to be willing to pay at least a 75% markup in order to generate the existence of a private unemployment insurance market, far exceeding willingness-to-pay estimates. I derive semi-parametric point estimates of this markup in excess of 300%.

In response to learning about future unemployment, I show that individuals decrease consumption and spouses are more likely to enter the labor market, even conditional on the later realization of unemployment. This suggests social insurance is valuable not only as insurance against the event of becoming unemployed, but also as insurance against learning about potentially becoming unemployed. While traditional formulas do not capture this value of social insurance, I provide several new methods to identify the ex-ante value of unemployment insurance. The methods exploit the impact of information revelation about future unemployment, in contrast to the impact of the event. The results suggest that, ex-ante before learning their own unemployment prospects, individuals would be willing to pay at least a 40-60% markup over an actuarially fair price for unemployment insurance.

*Harvard University, nhendren@fas.harvard.edu. I am very grateful to Daron Acemoglu, Raj Chetty, Amy Finkelstein, Jon Gruber, and Rob Townsend for their comments. I also thank Trevor Bakker, Augustin Bergeron, Lizi Chen, Frina Lin, Jeremy Majerovitz, Jimmy Narang, and Nina Roussille for helpful research assistance. An early version of the first half of this paper appeared as the second chapter in my MIT PhD thesis. Support from NSF Graduate Research Fellowship and the NBER Health and Aging Fellowship, under the National Institute of Aging Grant Number T32-AG000186 is gratefully acknowledged.

1 Introduction

The risk of losing one's job is arguably one of the most salient risks faced by working-age individuals. Job loss leads to drops in consumption and significant welfare losses.¹ Millions of people hold life insurance, health insurance, liability insurance, and many other insurance policies.² Why isn't there an analogous thriving market for insurance against losing one's job?³

Relatedly, the government is heavily involved in providing unemployment insurance (UI) benefits, and there is a growing literature characterizing the optimal amount of these benefits.⁴ Yet it is not clear what market failures, if any, provide a rationale for government intervention. If there is a welfare improvement from additional UI, why can't private firms provide such benefits? If knowledge about future unemployment creates a wedge between what the government and private markets can do, does this micro-foundation alter the calculus of the optimal amount of UI benefits?

This paper provides empirical evidence that unemployment or job loss insurance contracts would be too adversely selected to deliver positive profit, at any price. This generates a wedge between what can be accomplished by private markets and what can be achieved through government intervention that can require people to pay for insurance. Moreover, this micro-foundation modifies the formulas characterizing the utilitarian-optimal unemployment insurance benefit level.

Part 1 I proceed in two parts. In the first part I develop the argument that private information prevents the existence of a private UI market. I begin by developing a theory for when a UI market can exist and use the model to derive the empirical estimands of interest. Individuals

¹See Gruber (1997), Browning and Crossley (2001), Aguiar and Hurst (2005), and Chetty (2008) among others.

²60% of people in the US have insurance against damaging their cell phones and 1.4 million pets have health insurance in North America (see <http://www.warrantyweek.com/archive/ww20131114.html> and <http://www.embracepetinsurance.com/pet-industry/pet-insurance/statistics>).

³In the past 10 years, 2 companies have attempted to sell private unemployment insurance. PayCheck Gaurdian attempted to sell policies from 2008-2009, but stopped selling in 2009 with industry consultants arguing "The potential set of policyholders are selecting against the insurance company, because they know their situation better than an insurance company might" (<http://www.nytimes.com/2009/08/08/your-money/08money.html>). More recently, IncomeAssure has partnered with states to offer top-up insurance to a 50% replacement rate for workers in some industries and occupations (<https://www.incomeassure.com>). Back-of-the envelope calculations suggest their markups exceed 500% over actuarially fair prices. Indeed, it has been criticized for not saliently noting in its sales process that the government provides the baseline 30-40% replacement rate, shrouding the true price of the insurance (e.g. http://www.mlive.com/jobs/index.ssf/2011/08/get_out_your_calculator_before_you_buy_p.html#).

⁴See, for example, Baily (1976); Gruber (1997); Chetty (2008); Landais (2015) among many others.

may have private information about their future unemployment prospects, and insurance may increase their likelihood of unemployment through a moral hazard problem. In this environment, a private market cannot exist unless someone is willing to pay the pooled cost of those with higher probabilities of unemployment in order to obtain a small amount of insurance.⁵ This pooled cost depends on the distribution of job loss probabilities but does not depend on the responsiveness of unemployment to UI benefits. The first dollar of insurance provide first-order welfare gains, whereas the behavioral response imposes a second order impact on the cost of insurance, a point recognized by Shavell (1979). So although the moral hazard elasticity is useful for characterizing optimal social insurance, it does not readily provide insight into why a private market does not exist. This suggests that the key empirical measure of interest to characterize market existence is the markup over actuarially fair insurance that individuals would have to be willing to pay to cover the cost of worse risks.

I identify lower bounds and point estimates for these markups by building on an approach developed in Hendren (2013b) that uses subjective probability elicitations from the Health and Retirement Survey. Individuals are asked “what is the percent chance (0-100) that you will lose your job in the next 12 months?”. I do not assume individuals necessarily report their true beliefs on these survey questions; rather, I combine the elicitations with ex-post information about whether the individual actually loses her job to infer properties of the distribution of private information in the population. Under this approach, individuals have private information if they are able to predict their future job loss conditional on the observable characteristics insurers would use to price the insurance contracts, such as industry, occupation, demographics, unemployment history, etc.

Across a wide range of specifications, I find individuals hold a significant amount of private information that is not captured by the large set of observable characteristics available in the HRS. I use the distribution of predicted values of unemployment given the elicitations to show that individuals would have to be willing to pay at least a 70% markup on average to cover the pooled cost of worse risks. The presence of this private information is consistent across subsamples: old and young, long and short job tenure, industries and occupations, regions of the country, age groups, and across time. Using a parametric measurement error model, I move from lower bounds to semi-parametric point estimates that suggest individuals would need to

⁵This generalizes the no-trade condition of Hendren (2013b) to allow for moral hazard.

pay markups in excess of 300% in order to start an insurance market. This is well in excess of existing measures of individuals' willingness to pay for UI (generally between 15-50%).⁶

In health-related insurance markets, it is common to exclude high-risk individuals and offer insurance only to lower risk individuals. Hendren (2013b) presents evidence that private information drives this pattern (there's one way to be healthy and many unobservable ways to be sick). To explore whether such an underwriting strategy would mitigate the potential adverse selection problem with a private UI market, I restrict the analysis to "good risks" who have long job tenures and no previous unemployment spells. This does not mitigate the impact of potential adverse selection – if anything, such subgroups have higher frictions imposed by private information. Colloquially, there appears to be a 'bad apple in every bunch' that knows s/he may lose his/her job. This imposes significant costs on the vast majority of the population that has little knowledge about this event and has a low chance of losing their job.

Part II If knowledge about future unemployment prevents the existence of private UI, how does this affect how one should analyze optimal government intervention? Part II of the paper considers the optimal setting of unemployment benefits for a utilitarian planner when individuals may learn to some extent about future unemployment before the event occurs. The canonical optimality condition equates the marginal-utility smoothing benefits of UI to its fiscal cost inclusive of fiscal externalities from increased unemployment duration or other behavioral responses. Crucially, the marginal utility smoothing benefits are often measured using the causal effect of unemployment on the marginal utility of consumption (Baily (1976); Chetty (2008); Landais (2015)).

But when individuals learn ex-ante about future unemployment prospects, the causal effect of unemployment on the individuals' marginal utilities is no longer sufficient for welfare analysis. If individuals learn ex-ante about their potential future job loss, they may take aversive action to mitigate its effects. I document that in response to learning about the potential of losing one's job, individuals lower their consumption and spouses are more likely to enter the labor force. As a result, the benefits of UI are not simply given by the causal effect of unemployment on the marginal utility of consumption. Intuitively, UI delivers insurance not only against the event of unemployment, but also against the *realization of information* about future unemployment.

⁶As shown in Appendix Figure II, this also exceeds markups identified in Hendren (2013b) for those with pre-existing conditions in LTC, Life, and Disability – settings where the private market does not exist.

Although the presence of knowledge about future unemployment renders the causal effect of unemployment on marginal utilities a potentially biased measure of the value of social insurance, it does open up new directions for valuing social insurance. One can instead analyze the impact of information about unemployment on marginal utilities, as opposed to the impact of unemployment itself. I implement several new approaches following this direction.

The first approach utilizes consumption data in the HRS to estimate how ex-post consumption varies with the ex-ante likelihood of becoming unemployed. I document individuals who expect to become unemployed have lower consumption, even conditional on subsequent unemployment status. I illustrate how to use these responses to construct lower bounds on the value of UI. For a coefficient of relative risk aversion of 2, it suggests individuals would be ex-ante willing to pay at least a 50% markup for unemployment insurance. Second, in response to learning about future job loss, spouses are more likely to enter the labor market. Comparing these responses to an extensive margin spousal labor supply semi-elasticity of 0.5, it suggests individuals would be willing to pay a 60% markup to obtain insurance against learning one would become unemployed. Both of these estimates can be compared to estimates of the markup that would be required to cover the fiscal externality of providing additional unemployment benefits (e.g. from increased duration), such as 0.50 (Chetty (2008)).

Finally, I document that consumption responds to learning about future unemployment (even before unemployment occurs). Such identification not straightforward because the HRS – which contains subjective probability elicitation – does not contain information on consumption at the same time the elicitation is provided. Instead, I develop a 2-sample IV strategy that exploits the evolution of beliefs prior to unemployment (measured in the HRS) and the evolution of consumption prior to unemployment in the Panel Study of Income Dynamics (PSID). In response to unemployment in period t , I show consumption drops in year $t - 2$ relative to $t - 1$ even amongst those who remain employed in both previous years. Scaling this consumption drop by the amount of information revealed in year $t - 1$ relative to $t - 2$ (measured in the HRS) delivers an estimate of the impact of beliefs about future unemployment on consumption. This reveals a 10% increase in beliefs about future unemployment leads to a 2.5% drop in consumption today. Scaling by a coefficient of relative risk aversion of 2, this approach suggests individuals are ex-ante willing to pay at least a 50% markup for unemployment insurance.

More generally, if knowledge about future adverse events is the micro-foundation for gov-

ernment intervention, it suggests the canonical methods of welfare analysis need to be modified to take into account how marginal utilities vary with the realization of information about the event. The direct causal effect of the event on marginal utilities is not sufficient for the ex-ante welfare benefits of insurance against that event. But exploiting the impact of the evolution of information about the event provides a potentially fruitful path for measuring its total value.

Related literature This paper is related to a growing strand of literature studying unemployment insurance and, in particular, the optimal government provision of it.⁷ I show private information provides a micro-foundation for the absence of a private UI market and potential rationale for government intervention even in a partial equilibrium setting. In doing so, the model yields a modified formula for optimal social insurance that accounts for behavioral responses to information about the future adverse events, not just the onset of the events.⁸

In this sense, the paper is related to a large literature documenting behavioral responses to knowledge about future adverse events, including precautionary savings.⁹ Most closely, this paper is related to the work of Stephens (2004) who illustrates that subjective probability elicitation in the HRS are predictive about future unemployment status, and Stephens (2001) who documents the consumption patterns around unemployment.¹⁰ The findings of the spousal labor supply response is also related to the large literature studying the “added worker hypothesis” whereby spouses enter the labor market in response to unemployment shocks.¹¹ Relative to this literature, my results suggest a portion of the impact of unemployment on spousal labor supply may occur in response to beliefs about future unemployment prior to its actual event.

This paper also contributes to the growing literature documenting the impact of private information on the workings of insurance markets. Relative to this literature that often focuses on estimating adverse selection of existing contracts, it provides support for the hypothesis that

⁷See, for example, Baily (1976); Acemoglu and Shimer (1999, 2000); Chetty (2006); Shimer and Werning (2007); Chetty (2008); Shimer and Werning (2008); Landais et al. (2010).

⁸The Baily-Chetty formula has been modified in many ways. The modification proposed in this paper is most closely related to that of Andrews and Miller (2013) who consider the case of preference heterogeneity and show that one must account for the covariance between risk aversion and the consumption drop. In contrast, the modification proposed here arises even under homogeneous preferences because of the endogenous choices that occur in response to information realizations.

⁹See, for example, Barceló and Villanueva (2010); Bloemen and Stancaelli (2005); Carroll et al. (2003); Carroll and Samwick (1998, 1997); Dynan (1993); Engen and Gruber (2001); Guariglia and Kim (2004); Guiso et al. (1992); Hubbard et al. (1994); Lusardi (1997, 1998); Stephens (2001).

¹⁰Stephens (2004) also utilizes a sample of food expenditure information in the HRS and does not find that the subjective probability elicitation are correlated with food expenditure conditional on unemployment.

¹¹See, for example, Lundberg (1985); Maloney (1991); Gruber and Cullen (1996).

previous literature has perhaps suffered from a “lamp-post” problem, as suggested by Einav et al. (2010). If private information prevents the existence of entire markets, it is difficult to identify its impact by studying existing contracts.

The rest of this paper proceeds as follows. Part I (Sections 2-4) illustrates the frictions imposed by private information on the workings of a private UI market. Section 2 presents the theoretical model and derives the estimands that characterize the frictions imposed by private information. Section 3 describes the data used; Section 4 estimates the frictions imposed by private information. Part II (Sections 5-6) considers the implications for government intervention. Section 5 presents a modified Baily-Chetty formula characterizing the optimal level of government benefits, illustrating how behavioral responses to information can affect this valuation. Section 6 provides estimates of the behavioral responses to information about unemployment on consumption and spousal labor supply, and quantifies these impacts on the value of social insurance. Section 7 concludes.

Part I: Private Information as a Barrier to Private UI Markets

2 Theory

I consider a theoretical model of unemployment risk. The goal of the model is to derive the estimands that will form the basis of the empirical work to characterize when a private market can exist (Part I) and to characterize the ex-ante (utilitarian) optimal level of social insurance (Part II). The model will contain both private information (which will create a wedge between what private markets and the government can achieve in terms of allocations) and moral hazard (which implies full insurance is not necessarily socially optimal).

2.1 Setup

There exists a unit mass of currently employed individuals indexed by an unobservable type $\theta \in \Theta$. While θ is unobserved, individuals have observable characteristics, X , that insurers could potentially use to price insurance contracts. Individuals face a potential of losing their job, which occurs with probability p . Individuals of type θ make a set of choices from some feasible choice set to maximize their utility. These choices include consumption in the event of being employed, consumption in the event of being unemployed, the probability of losing their job, p , and a range of other actions, a , that can include items like future consumption, labor effort, and spousal labor supply. These choices are made subject to a choice set $\{c_e, c_u, p, a\} \in \Omega(\theta)$ which may vary across types. Individuals have increasing and concave utility functions over consumption in the state of being unemployed, $u(c)$, and the state of being employed, $v(c)$, which may differ to allow for state dependent utility, and a (separable) utility function $\Psi(1-p, a; \theta)$ over other choices, a , and the probability of being employed, p .

In addition, there may or may not exist an insurance policy that pays b in the event of being unemployed at a premium of τ paid in the event of being employed. The aggregate utility of such a policy is given by

$$U(\tau, b; \theta) = \max_{\{c_e, c_u, p, a\} \in \Omega(\theta)} (1-p)v(c_e - \tau) + pu(c_u + b) - \Psi(1-p, a; \theta) \quad (1)$$

where $c_e - \tau$ is the net consumption in the event of employment and $c_u + b$ is the net consumption in the event of unemployment. Individuals choose c_e and c_u after knowing b and τ , so that one

could equivalently think of the individual as choosing consumption.¹² The main goal of Part I is to ask when private markets can profitably sell policies with non-zero values of b and τ . Part II will consider the socially optimal choice of b and τ .

There are two key frictions in the model. First, individuals have private information about their types, θ , and in particular their probability of becoming unemployed, $p(\theta)$. Second, individuals are able to potentially choose their probability of becoming unemployed, and hence there is a moral hazard problem.¹³ To see this, consider the case when $\Psi(1-p, a; \theta)$ is convex in $1-p$ so that the choice of p by type θ satisfies the first order condition:

$$v(c_e(\theta) - \tau) - u(c_u(\theta) + b) = \Psi'(1-p(\theta), a(\theta); \theta) \quad (2)$$

where $\Psi'(1-p, a(\theta); \theta)$ denotes the first derivative of Ψ with respect to $1-p$, evaluated at the individual's optimal allocation.¹⁴ Intuitively, the marginal cost of effort to avoid unemployment is equated to the benefit, given by the difference in utilities between employment and unemployment. Note that different types, θ , may have different underlying probabilities, $p(\theta)$, that satisfy equation (2).

2.2 Allocations Implementable through Private Markets

When can a private market profitably sell a private insurance policy, (b, τ) ? To start, consider a policy that provides a small payment, db , in the event of being unemployed and is financed with a small payment in the event of being employed, $d\tau$, offered to those with observable characteristics X . By the envelope theorem, the utility impact of buying such a policy will be given by

$$dU = -(1-p(\theta))v'(c_e(\theta))d\tau + p(\theta)u'(c_u(\theta))db$$

which will be positive if and only if

$$\frac{p(\theta)u'(c_u(\theta))}{(1-p(\theta))v'(c_e(\theta))} \geq \frac{d\tau}{db} \quad (3)$$

The LHS of equation (3) is a type θ 's willingness to pay (i.e. marginal rate of substitution) to move resources from the event of being employed to the event of being unemployed.¹⁵ The RHS

¹²Modeling the insurance contract in this manner removes the need to explicitly model the set of constraints, $\Omega(\theta)$.

¹³For simplicity, I assume the functions v and u are common to all types; this will be useful in the empirical application of the welfare analysis in Part II, but is not critical for Part I.

¹⁴Without loss of generality, one can suitably define Ψ so that the first order condition holds even if $p \in \{0, 1\}$.

¹⁵Note that, because of the envelope theorem, the individual's valuation of this small insurance policy is independent of any behavioral response. While these behavioral responses may impose externalities on the insurer

of equation (3), $\frac{d\tau}{db}$, is the cost per dollar of benefits of the hypothetical policy.

Let $\bar{\Theta} \left(\frac{d\tau}{db} \right)$ denote the set of all individuals, θ , who prefer to purchase the additional insurance at price $\frac{d\tau}{db}$ (i.e. those satisfying equation (3)) who have observable characteristics X . An insurer's profit from a type θ is given by $(1 - p(\theta)) \tau - p(\theta) b$. Hence, the insurer's marginal profit from trying to sell a policy with price $\frac{d\tau}{db}$ is given by

$$d\Pi = \underbrace{E \left[1 - p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right]}_{\text{Premiums Collected}} d\tau - \underbrace{E \left[p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right]}_{\text{Benefits Paid}} db - \underbrace{\left(dE \left[p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right] \right)}_{\text{Moral Hazard}} (\tau + b)$$

where the first term is the amount of premiums collected, the second term are the benefits paid out, and the third term is the impact of offering additional insurance on the cost of providing the baseline amount of insurance. Additional insurance may increase the cost through increased probability of unemployment, $dE[p(\theta)] > 0$.¹⁶ But, for the first dollar of insurance, the moral hazard cost to the insurer is zero because $\tau = b = 0$. This insight, initially noted by Shavell (1979), suggests moral hazard does not affect whether insurers' first dollar of insurance is profitable – a result akin to the logic that deadweight loss varies with the square of the tax rate.

The first dollar of insurance will be profitable if and only if

$$\frac{d\tau}{db} \geq \frac{E \left[p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right]}{E \left[1 - p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right]} \quad (4)$$

If inequality (4) does not hold for any possible price, $\frac{d\tau}{db}$, then providing private insurance will not be profitable at any price. The market will unravel a la Akerlof (1970).

A couple simplifications on the model environment facilitate a clearer expression of inequality (4). First, I assume that the mapping from types, θ , to probabilities, $p(\theta)$, is one-to-one so that there is no variation in marginal utilities conditional on $p(\theta)$. Hence a type θ 's willingness to pay for an additional unit of insurance (LHS of equation (3)) will not vary conditional on p and can be written $\frac{p}{1-p} \frac{u'(c_u(p))}{v'(c_e(p))}$ without loss of generality. Second, I assume that this willingness to pay, $\frac{p}{1-p} \frac{u'(c_u(p))}{v'(c_e(p))}$, is increasing in p – that is, those with higher probabilities are willing to pay more for insurance that pays when unemployed. Under these assumptions, the adverse selection will take a particular threshold form: the set of people who would be attracted to a contract for which type $p(\theta)$ is indifferent will be the set of higher risks whose probabilities

or government, they do not affect the individuals' willingness to pay.

¹⁶To incorporate observable characteristics, one should think of the expectations as drawing from the distribution of θ conditional on a particular observable characteristic, X .

exceed $p(\theta)$. Notationally, let P denote the random variable corresponding to the distribution of probabilities chosen in the population in the status quo world without a private unemployment insurance market, $b = \tau = 0$.¹⁷ Equation (4) can be re-written as:

$$\frac{u'(c_u(p))}{v'(c_e(p))} \leq T(p) \quad \forall p \quad (5)$$

where $T(p)$ is given by

$$T(p) = \frac{E[P|P \geq p]}{E[1 - P|P \geq p]} \frac{1 - p}{p}$$

which is the pooled cost of worse risks, termed the ‘‘pooled price ratio’’ in Hendren (2013b). The market can exist only if there exists someone who is willing to pay the markup imposed by the presence of higher risk types adversely selecting her contract. Here, $\frac{u'(c_u(p))}{v'(c_e(p))} - 1$ is the markup individual p would be willing to pay and $T(p) - 1$ is the markup that would be imposed by the presence of risks $P \geq p$ adversely selecting the contract. This suggests the pooled price ratio, $T(p)$, is the fundamental empirical magnitude desired for understanding the frictions imposed by private information.

In the case when there are two types θ with different willingnesses to pay but the same probability of unemployment, types do not map 1-1 into $p(\theta)$, and equation (4) does not summarize the no trade condition. However, Appendix A.1 shows that there exists a mapping, $f(p)$, from a subset of $[0, 1]$ into the type space, Θ , such that the no trade condition reduces to testing

$$\frac{u'(c_u(f(p)))}{u'(c_e(f(p)))} \leq T(p) \quad \forall p \quad (6)$$

Hence, the pooled price ratio continues to be a key measure for the frictions imposed by private information even in the presence of multi-dimensional heterogeneity.¹⁸

Minimum and average $T(p)$ What statistics of $T(p)$ are desired for estimation? The no trade condition in equation (5) must hold for all p . Absent particular knowledge of how the

¹⁷In other words, the random variable P is simply the random variable generated by the choices of probabilities, $p(\theta)$, in the population.

¹⁸Appendix A discusses the generality of the no trade condition. Appendix A.3 illustrates that while in principle the no trade condition does not rule out non-marginal insurance contracts (i.e. b and $\tau > 0$), in general a monopolist firm’s profits will be concave in the size of the contract; hence the no trade condition also rules out larger contracts. Appendix A.2 also discusses the ability of the firm to potentially offer menus of insurance contracts instead of a single contract to screen workers. Hendren (2013b) considers this more general case with menus in a model without moral hazard and shows that when the no trade condition holds, pooling delivers weakly higher profit than a separating contract. In Appendix A.2, I show that a version of the present model without the multi-dimensional heterogeneity can be nested into that model. I also discuss potential theoretical directions for future work studying the robustness of the no trade condition in equations (5) and (6).

willingness to pay for UI varies across p , it is natural to estimate the minimum pooled price ratio, $\inf T(p)$, as in Hendren (2013b). If no one is willing to pay this minimum pooled price ratio, then the market cannot exist.¹⁹

But, by taking the minimum one implicitly assumes that an insurer trying to start up a market would be able to a priori identify the best possible price that would minimize the markup imposed by adverse selection. In contrast, if insurers do not know exactly how best to price the insurance (e.g. because there is no market from which to learn the distribution of types), the price of adverse selection imposed on a potential market entrant could be higher and depend on other properties of the pooled price ratio. This can motivate the average pooled price ratio, $E[T(p)]$, as a complementary statistic for studying the degree of potential adverse selection.

To see this, suppose an insurer seeks to start an insurance market by randomly drawing an individual from the population and, perhaps through some market research, learns exactly how much this individual is willing to pay. Let's say this person has a probability p of becoming unemployed and for simplicity assume the mapping from types to p is one-to-one. The insurer offers a contract that collects \$1 in the event of being employed and pays an amount in the unemployed state that makes the individual perfectly indifferent to the policy. Then the insurer tries to sell this policy to the marketplace; clearly, all risks $P \geq p$ will choose to purchase the policy as well. Therefore, the profit per dollar of revenue will be

$$r(p) = \frac{u'(c_u(p))}{v'(c_e(p))} - T(p)$$

So, if the original individual was selected at random from the population, the expected profit per dollar would be positive if and only if

$$E \left[\frac{u'(c_u(p))}{v'(c_e(p))} \right] \geq E[T(P)] \tag{7}$$

If the insurer is randomly choosing contracts to try to sell, it is not the minimum pooled price ratio that determines profitability. Rather, on average, individuals would have to be willing to pay the pooled price ratio, $E[T(P)]$. In this sense, the average pooled price ratio provides guidance on the frictions imposed on a potential insurance company entrant that would attempt

¹⁹Although not a necessary condition, the no trade condition will hold if

$$\sup_{p \in [0,1]} \frac{u'(c_u(f(p)))}{v'(c_e(f(p)))} \leq \inf_{p \in [0,1]} T(p)$$

so that absent particular knowledge about how the willingness to pay varies across p , the minimum pooled price ratio provides guidance into the frictions imposed by private information.

to set up a market through experimentation. From a more practical standpoint, Section 4 will illustrate that one can construct lower bounds on $E[T(p)]$ under weaker assumptions than are required to estimate the minimum pooled price ratio. Hence, it will be useful to have in mind the theoretical relationship between $E[T(p)]$ and the barriers to trade imposed by private information.

3 Data

The model motivates a particular measure of private information: how much of a markup over actuarially-fairly priced insurance would individuals have to be willing to pay in order make it profitable for an insurance company to sell insurance? To answer this, I build on the approach of Hendren (2013b) which uses subjective probability elicitations as noisy and potentially biased measures of true beliefs to identify and quantify private information.²⁰

I use data from the Health and Retirement Study (HRS), a panel survey spanning years 1992-2013. The HRS samples individuals generally over 55 and their spouses (included regardless of age).²¹ To identify the frictions imposed by private information on the workings of a private unemployment insurance market, the survey contains three key sets of variables: a subjective probability elicitation about future unemployment, its corresponding indicator for whether or not unemployment occurs, and a set of public information insurance companies could use to price the insurance contracts.

Subjective probability elicitations. The survey asks respondents: what is the percent chance (0-100) that you will lose your job in the next 12 months? I denote these free-responses by Z . Figure I presents the histogram of the subjective probability elicitations. As has been noted in previous literature (Gan et al. (2005)), these responses tend to concentrate on focal point values, especially zero.²² This motivates the approach of not assuming these are true beliefs, but rather that they may be noisy and biased measures of true beliefs.

²⁰While standard approaches to identifying private information search for correlations between the degree of insurance purchase and subsequent insured losses or claims (Chiappori and Salanié (2000); Finkelstein and Poterba (2002, 2004); Einav et al. (2010)), such revealed-preference methods of identifying private information are not well-suited to identify it in cases when private information has led to an absence of a private market and limits individuals' choice sets.

²¹Despite its focus on an older set of cohorts, the HRS is a natural dataset choice because it contains information on unemployment, consumption, a wide range of observable characteristics insurers use in other markets to price policies, and, most importantly, subjective probability elicitations about future unemployment.

²²Taken literally, a response of zero or 100 implies an infinite willingness to pay for certain financial contracts, which contrasts with observed behavior.

Incidence of Job Loss. Corresponding to the elicitation, the panel nature of the survey facilitates the construction of whether or not the individual will involuntarily lose their job in the subsequent 12 months from the survey, denoted U . In particular, the subsequent wave asks individuals whether they are working at the same job as the previous wave (~ 2 years prior). If not, respondents are asked when and why they left their job. In particular, they ask whether they lost their job, left voluntarily/quit, or retired. To most closely align with the definition of unemployment and the subjective probability elicitation, I define becoming unemployed as involuntarily losing one's job in the subsequent 12 months following the previous survey date, and I exclude voluntary quits and retirement. As a result, the empirical work will estimate the frictions imposed by private information on a hypothetical insurance market that pays \$1 in the event the individual involuntarily loses his/her job in the subsequent 12 months.

I consider robustness analyses to other definitions of job loss. I construct a measure of job loss in the 6-12 months following the survey. This removes cases where the individuals knew about an immediately impending job loss that could potentially be circumvented by an insurer imposing a waiting period on the insurance policy. I also construct measures of job loss in the 6-24 month window, and measures of whether the individual is unemployed in the subsequent survey round (roughly 24 months after the previous survey). Finally, I construct measures of job loss that are the product of these indicators with an indicator for receiving positive government unemployment insurance benefits in between survey waves, thereby restricting to the set of job losses that led to a government UI claim.

Public Information. Estimating private information requires specifying the set of observable information insurers could use to price insurance policies. The data contain a very rich set of observable characteristics that well-approximate variables used by insurance companies in disability, long-term care, and life insurance (Finkelstein and McGarry (2006); He (2009); Hendren (2013b)) and also contain a variety of variables well-suited for controlling for the observable risk of job loss. The baseline specification includes a set of these job characteristics including job industry categories, job occupation categories, log wage, log wage squared, job tenure, and job tenure squared, along with a set of demographic characteristics (census division dummies, gender dummies, age, age squared, and year dummies).²³

²³This set is generally larger than the set of information previously used by insurance companies who have tried to sell unemployment insurance. Income Assure, the latest attempt to provide private unemployment benefits, prices policies using a coarse industry classification, geographical location (state of residence), and wages.

I also assess robustness to additional health status controls that include indicators for a range of doctor-diagnosed medical conditions (diabetes, a doctor-diagnosed psychological condition, heart attack, stroke, lung disease, cancer, high blood pressure, and arthritis) and linear controls for bmi.²⁴ I also consider specifications that condition on lagged unemployment incidence, and also to a less comprehensive set of controls such as just age and gender. By changing the set of observable characteristics, the empirical results allow one to understand how the potential for adverse selection varies with the underwriting strategy of the potential insurer.

Sample. I begin with a sample of everyone under 65 currently holding a job who is asked the subjective probability elicitation question, Z . I keep only those respondents who have non-missing job loss responses in the subsequent wave, U , and those with non-missing observable characteristics, X . I exclude the self-employed and those employed in the military.

Table I presents the summary statistics of the samples used in the paper. There are 26,640 observations in the sample, which correspond to 3,467 unique households. The average age is 56 and roughly 40% of the sample is male. Because the HRS primarily focuses on an older population, the mean age in the sample is 56. Although the average age is quite high, I present evidence below that the patterns are actually quite stable across the age ranges I can observe. Mean yearly wages are around \$36,000 in the baseline sample and average job tenure is 12.7 years.

In the subsequent 12 months from the survey, 3.1% of the sample reports losing their job involuntarily. In contrast, the mean subjective probability elicitation is 15.7%, indicating a significant bias in elicitation on average, and a wide standard deviation of 0.173. This is arguably a well-known artifact of the non-classical measurement error process inherent in subjective elicitation. Elicitations are naturally bounded between 0 and 1. Hence, for low probability events, there is a natural tendency for measurement error in elicitation to lead to an upward bias in elicitation. This provides further rationale for treating these elicitation as noisy and potentially biased measures of true beliefs, as is maintained throughout the empirical analyses below.

²⁴As shown in Panel 2 of Table 1, 22,831 observations of the 26,640 baseline observations report non-missing values for these health variables.

4 Empirical Evidence of Private Information

4.1 Presence of Private Information and Lower Bounds on $E [T (P)]$

I begin by documenting the relationship between the subjective probability elicitation, Z , and subsequent unemployment, U , conditional on the observable demographic and job characteristics, X . To illustrate the predictive content in the elicitation, Figure II (Panel A) bins the elicitation into 5 groups and presents the coefficients on these indicators in a regression of U on these bin dummies and the observable controls, X . The vertical axis reports these coefficients. The horizontal axis presents the mean value of the elicitation in each bin. The figure displays a clear increasing pattern: those with higher subjective probability elicitation are more likely to lose their job, conditional on demographics and job characteristics.

While Figure II (Panel A) presents evidence that individuals have knowledge about their future unemployment prospects, it does not provide information on the pooled price ratio, $T (p)$. To relate the predictive content in the elicitation to the frictions it can impose on the workings of an insurance market, one can proceed in several steps. First, consider the predicted values

$$P_Z = \Pr \{U|X, Z\}$$

Under a couple of natural assumptions, Hendren (2013b) shows that the distribution of predicted values forms a distributional lower bound on the distribution of true beliefs.

Remark 1. (Hendren (2013b)) Suppose (a) elicitation contain no more information about U than does P : $\Pr \{U|X, Z, P\} = \Pr \{U|X, P\}$ and (b) true beliefs are unbiased $\Pr \{U|X, P\} = P$. Then true beliefs are a mean-preserving spread of the distribution of predicted values:

$$E [P|X, Z] = P_Z$$

The predicted values, P_Z , present a distributional lower bound on the true distribution of beliefs in the population – the true distribution is more dispersed than the observed distribution of predicted values.

Figure II, Panel B constructs the distribution of the predicted values of $P_Z - \Pr \{U|X\}$. To construct this figure, I use a probit specification²⁵ in X and Z that includes a second order polynomial in Z to capture the potential nonlinearities, such as the moderately convex relationship

²⁵Results are similar using a linear specification (as shown in Appendix Table I), but since the mean probability of becoming unemployed is very close to zero (3.1%) the probit specification has a better fit since the specification is not fully saturated in X and Z .

illustrated in Figure II, and also indicators for $Z = 0$, $Z = 0.5$, and $Z = 1$ to capture focal point responses illustrated in Figure I. This produces the predicted values, P_Z . To construct $\Pr\{U|X\}$, I run the same specification but exclude the Z variables.

If individuals had no private information, the distribution in Figure II, Panel B, would be statistically identical to a point mass at 0. While there is a large mass of risks with very low predicted probabilities of unemployment, there is an “upper tail” of predicted probabilities lying above the mass of low-risks. In order to start a profitable insurance market, the mass of low-risks would need to be willing to pay a large enough markup to cover the costs of these higher risks.

To understand how much of a markup these higher risks may impose, I extend some results initially derived in Hendren (2013b). For expositional simplicity, consider a particular observable characteristic, $X = x$ and define $m(p) = E[P - p|P \geq p]$ to be the mean residual life function of the distribution P for those with $X = x$. Intuitively, $m(p)$ asks “how much worse are the worse risks than p ?” Of course, $m(p)$ is not observed without observing the true distribution of P . But, one can construct a sample analogue of $m(p)$ using the distribution of predicted values, P_Z :

$$m_Z(p) = E[P_Z - p|P_Z \geq p]$$

Hendren (2013b) shows that the average value of $m_Z(P_Z)$ generates a lower bound on the average value of $m(P)$:

$$E[m_Z(P_Z)] \leq E[m(P)] \tag{8}$$

Therefore, one can use the observed distribution of the predicted values, P_Z , to generate a lower bound on the extent to which worse risks would impose costs on lower risks if a market were started. With heterogeneous values of X , equation (8) will hold for each value of X (and thus it will hold in expectation as well).

Proposition 1 goes one step further by using the distribution of predicted values to form a lower bound on the average pooled price ratio, $E[T(P)]$.

Proposition 1. *Suppose (a) elicitation contain no more information about U than does P : $\Pr\{U|X, Z, P\} = \Pr\{U|X, P\}$ and (b) true beliefs are unbiased $\Pr\{U|X, P\} = P$. Then,*

$$E[T(P)] - 1 \geq E[T_Z(P_Z) - 1] \tag{9}$$

where

$$T_Z(P_Z) = 1 + \frac{m(P_Z)}{\Pr\{U\}}$$

Proof. (See Appendix B) The proof extends results in Hendren (2013b) by applying Jensen’s inequality to $T(P)$. □

The extent to which the average pooled price ratio, $E[T(p)]$, exceeds 1 is bounded below by the ratio of $E[m_Z(P_Z)]$ and $\Pr\{U\}$. Intuitively, $E[m_Z(P_Z)]$ provides a lower bound on the extent to which risks have higher probabilities than p , so that the ratio relative to the mean probability, $\Pr\{U\}$, provides a lower bound on the average pooled cost that one would have to pay to cover the cost of the higher risks.

Table II presents the results.²⁶ For the baseline specification with demographic and job characteristic controls, the average markup imposed by the presence of worse risks is at least 76.82% (i.e. $E[T(P)] \geq 1.7682$). Adding health controls changes this slightly to 72%; dropping the job characteristic controls increases this slightly to 80%. The presence of such markups impose significant barriers to the existence of a private insurance market for UI.

Figure III (Panel A) presents the results graphically for the specifications with alternative control variables (X) against the psuedo-R squared of the model for $\Pr\{U|X, Z\}$. As one can see, including job characteristics significantly increases the predictive power of the model, but it does not meaningfully reduce the barrier to trade imposed by private information. Intuitively, the additional job characteristics controls help better predict unemployment entry rates across, say, industry and occupation groups; but it does not remove the thick upper tail illustrated in Figure II, Panel B.²⁷

Adding further controls does not appear to significantly modify the frictions imposed by private information, nor does it significantly alter the R-squared of the model. Controlling for health information does not meaningfully change the estimates, nor does adding additional con-

²⁶As in Hendren (2013b), the construction of $E[T_Z(P_Z)]$ and $E[m_Z(P_Z)]$ is all performed by conditioning on X . To partial out the predictive content in the observable characteristics, I first construct the distribution of residuals, $P_Z - \Pr\{U|X\}$. I then construct $m_Z(p)$ for each value of X as the average value of $P_Z - \Pr\{U|X\}$ above $p + \Pr\{U|X\}$ for those with observable characteristics X . In principle, one could estimate this separately for each X ; but this would require observing a rich set of observations with different values of Z for that given X . In practice, I follow Hendren (2013b) and specify a partition of the space of observables, ζ_j , for which I assume the distribution of $P_Z - \Pr\{U|X\}$ is the same for all $X \in \zeta_j$. This allows the mean of P_Z to vary richly with X , but allows a more precise estimate of the shape by aggregating across values of $X \in \zeta_j$. In principle, one could choose the finest partition, $\zeta_j = \{X_j\}$ for all possible values of $X = X_j$. However, there is insufficient statistical power to identify the entire distribution of P_Z at each specific value of X . For the baseline specification, I use an aggregation partition of 5 year age bins by gender. Appendix Table I (Columns (3)-(5)) documents the robustness of the results to alternative aggregation partitions.

²⁷Appendix Table I explores robustness to various specifications, including linear versus probit error structures, alternative aggregation windows for constructing $E[m_Z(P_Z)]$, and alternative polynomials for Z . All estimates are quite similar to the baseline and yield lower bounds of $E[T_Z(P_Z)] - 1$ of around 70%.

trols for their work history such as controls for indicators for being employed in the previous two survey waves, as indicated by the “Demo, Job, History” specification. Conversely, dropping the demographic variables such as region, year, gender etc and solely using age and age squared leads to a similar magnitude of private information relative to the baseline specification. Intuitively, the friction imposed by private information is not so much driven by variation in mean odds of becoming unemployed, but rather the thick upper tail of personally-specific knowledge that an individual may have that he or she has a particular chance at losing his or her job.

To illustrate the difficulty faced by a potential insurer in removing the information asymmetry, Figure III, Panel B adds individual fixed effects to a linear specification for $\Pr\{U|X, Z\}$.²⁸ Of course, such fixed effects would be impossible for an insurer to use – an econometrician can view the fixed effects as nuisance parameters that drop out in a linear fixed effects model; in contrast, an insurer must view them as a key input into their pricing policy.²⁹ Adding these fixed effects significantly increases the R^2 of the model, but individuals would still on average have to be willing to pay at least a 40% markup to cover the pooled cost of worse risks. Relatedly, while the autocorrelation in Z across waves is around 0.25, there exists significant predictive content within person, which is consistent with the individual’s elicitation containing largely personal and time-varying knowledge about future job loss.

Population Heterogeneity While on average the markups individuals would have to be willing to pay to start a private UI market are high, a natural question to ask is whether there are certain subsets of the population for which these would be lower. Columns (4)-(9) of Table II and Figure III, Panels C-F present the estimates of $E[T(P)] - 1$ for various subsamples of the data. Panels C and D of Figure III illustrate the presence of a large amount of private information across all industries and occupational subgroups, with lower bounds on $E[T(P)]$ that all exceed 50%.

The presence of significant amounts of private information about future job loss also spans the age spectrum in the data (45-65), as shown in Columns (4)-(5) of Table III and Panel E of Figure

²⁸I use the linear specification so that the residuals, $P_Z - \Pr\{U|X\}$ are well identified and do not suffer bias from the inability to consistently estimate the nuisance parameters. Appendix Table I, Column (2) illustrates that the baseline value for $E[T_Z(P_Z)] - 1$ is 0.6802 (s.e. 0.051) when using the linear specification for $\Pr\{U|X, Z\}$ as opposed to the baseline value of 0.7687 using the probit specification. Hence, a small amount of the attenuation illustrated in Figure III, Panel B (where the fixed effects specification yields 0.40) for the fixed effects estimates relative to the baseline is driven by the specification change from probit to linear.

²⁹Moreover, the econometrician is able to construct these fixed effects ex-post (after observing U realizations for the individual over many years), whereas an insurer would generally attempt to construct this ex-ante.

III. Columns (6)-(7) of Table II splits the sample by below and above-median wage workers and continues to find consistent evidence of private information across these subgroups, with lower bounds of 65% and 95%. Appendix Figure I, Panels A and B illustrates the consistency of private information across all years and across all census divisions in the U.S.

One underwriting strategy that has been common in other insurance markets is to limit the insurance market to “good risks”. For example, health-related insurance markets generally exclude those with pre-existing conditions. Hendren (2013b) shows this is consistent with those risks having private information but healthy individuals not.³⁰

Figure III, Panel F asks whether a similar underwriting strategy could help open up an unemployment insurance market for those with a low chance of losing their job. The figure plots the estimated $E[T_Z(P_Z)] - 1$ for varying subsamples with different job tenure and work histories. In contrast to the idea that restricting to good risks would help open up an insurance market, the figures illustrate if anything the opposite pattern: better risk populations have higher markups. Indeed, for those with greater than 5 years of job tenure, the data suggest a lower bound of 110% despite having a less than 2% chance of losing their job in the subsequent 12 months.

Loosely, the data is consistent with there always being at least one bad apple in every bunch that knows s/he has a decent chance of losing his/her job. This presents an especially high burden on a sample that have very low probabilities of unemployment, leading to higher implicit markups for these groups and preventing insurers from opening up markets to those who, based on observables, seem like especially good risks.

Alternative outcomes and waiting periods The results suggest high markups imposed by private information on a hypothetical insurance market that pays \$1 in the event of becoming unemployed in the subsequent 12 months. One alternative market – which would be consistent with insurance policies in other contexts – would be to impose waiting periods of, for example, 6 months before the insurance goes into effect. Indeed, if the private information is primarily about knowing that one will lose their job next week, then excluding next week from the insurance contract payouts could remove the informational asymmetry.

³⁰Loosely, those results suggest that there’s one way to be healthy, but many unobservable ways to be sick. This pattern prevents the existence of insurance markets for those with pre-existing conditions, but the ability of insurers to limit such risks from risk pools allows for insurance markets for the healthy that are less afflicted by problems of private information.

In Appendix Table I, I consider an alternative definition of U that excludes those who become unemployed in the first 6 months after the survey. I continue to use the same elicitation, Z , in the construction of the distribution of predicted values. This is appropriate because Z can still satisfy the assumptions in Remark 1 for the alternative measure of U ; but is likely to be a noisier measure of the individual’s true beliefs about losing his or her job in the 6-12 months after the survey, as opposed to the 0-12 months after the survey, as is prompted in the elicitation. Hence, one might expect lower values for $E[T_Z(P_Z)]$ because of this additional measurement error, but it remains a lower bound for the true markup that would be imposed by the presence of private information for an insurance contract that paid db in the event of unemployment with a 6-month waiting period.³¹

In practice, the results imply a lower bound on $E[T(P)] - 1$ of 0.579 ($p < 0.001$) for a market that imposes a 6-month waiting period. This suggests the frictions imposed by private information cannot be removed through the imposition of waiting periods.

Another strategy could be to require individuals to also file for unemployment insurance with the government.³² Such a practice could impose higher take-up hurdles and also help mitigate claims from job loss events that don’t lead to significant periods of unemployment. To assess the potential barriers to trade imposed by private information in such a market, I construct an outcome that is the interaction of unemployment with whether or not the individual receives government UI benefits. Appendix Figure I, Panel C plots the estimated lower bounds, $E[T_Z(P_Z)] - 1$, for such a hypothetical market. Restricting to government UI for a 0-12 month contract has a lower bound on the average markup of roughly 95%. The figure also illustrates that the implicit markups remain high for other potential timelines, such as 0-24 and 6-24 month payout windows. Restricting insurance payouts to cases in which the individuals filed government UI benefits would not appear to significantly reduce the barriers to trade imposed by private information.

Overall, the results document significant lower bounds on the average markups individuals would have to be willing to pay in order to cover the pooled cost of worse risks. They generally

³¹I also abstract from the ability of an individual to change the timing of their unemployment. Such claim timing could impose additional adverse selection costs. In principle, if such timing responses are costly to the worker, they would be a behavioral response that would not affect the insurer’s costs for the first dollar of insurance when $b = \tau = 0$. But, this could be an additional cost factor with non-marginal contracts, as has been noted in other market contexts such as dental insurance (Cabral (2013)).

³²Indeed, this is part of the strategy taken by the most recent attempt at providing unemployment insurance by Income Assure.

exceed 50% across a wide set of specifications, subsamples, and controls for observable characteristics. Moreover, these lower bounds are derived solely using the assumptions outlined in Remark 1 that allow the elicitation to be noisy and potentially biased measures of true beliefs. The next section adds additional assumptions about the nature of the measurement error in the elicitation that allows one to move from a lower bound on $E[T(P)]$ to point estimates for $T(p)$ and its minimum, $\inf T(p)$.

4.2 Quantification of $\inf T(p)$

To generate a point estimate for the pooled price ratio, one requires an estimate of the distribution of beliefs, P . To obtain this, I follow Hendren (2013b) by making additional assumptions about the distribution of measurement error in the elicitation. Note that the observed density (p.d.f./p.m.f.) of Z and U can be written as

$$f_{Z,U}(Z, U|X) = \int_0^1 p^U (1-p)^{1-U} f_{Z|P,X}(Z|P=p, X) f_P(p|X) dp$$

where $f_{Z|P,X}$ is the distribution of elicitation given true beliefs (i.e. elicitation error) and f_P is the distribution of true beliefs in the population (which can be used to construct $T(p)$ at each p). This is obtained by first taking the conditional expectation with respect to p and then using the assumption that $\Pr\{U|Z, X, P\} = P$.

To estimate the distribution of beliefs, f_P , I assume that the distribution of elicitation error, $f_{Z|P}(Z|P)$ can be represented by a low-dimensional vector of parameters; I then estimate these parameters along with a flexible specification for the distribution of true beliefs, $f_P(p|X)$.

I follow Hendren (2013b) by assuming that $Z = P + \epsilon$, where ϵ has the following structure. With probability λ , individuals report a noisy measure of their true belief P that is drawn from a $[0, 1]$ -censored normal distribution with mean $P + \alpha(X)$ and variance σ^2 . With this specification, $\alpha(X)$ reflects potential bias in elicitation and σ represents the noise. While this allows for general measurement error in the elicitation, it does not produce the strong focal point concentrations shown in Figure 1 and documented in existing work (Gan et al. (2005)). To capture these, I assume that with probability $1 - \lambda$ individuals take their noisy report with the same bias $\alpha(X)$ and variance σ^2 , but censor it into a focal point at 0, 50, or 100. In particular, if their elicitation would have been below κ , they report zero. If it would have been between κ and $1 - \kappa$, they report 50; and if it would have been above $1 - \kappa$, they report 1. Hence, I estimate

four elicitation error parameters: $(\sigma, \lambda, \kappa, \alpha(X))$ that capture the patterns of noise and bias in the relationship between true beliefs, P , and the elicitation reported on the surveys, Z .³³

Ideally, one would flexibly estimate the distribution of P given X at each possible value of X . This would enable separate estimates of the minimum pooled price ratio for each value of X . However, the dimensionality of X prevents this in practice. Instead, I again follow Hendren (2013b) and adopt an index assumption on the cumulative distribution of beliefs, $F(p|X) = \int_0^p f_P(\tilde{p}|X) d\tilde{p}$,

$$F(p|X) = \tilde{F}(p|\Pr\{U|X\}) \quad (10)$$

where I assume $\tilde{F}(p|q)$ is continuous in q (where $q \in \{0, 1\}$ corresponds to the level of $\Pr\{U|X\}$). This assumes that the distribution of private information is the same for two observable values, X and X' , that have the same observable unemployment probability, $\Pr\{U|X\} = \Pr\{U|X'\}$. Although one could perform different dimension reduction techniques, controlling for $\Pr\{U|X\}$ is particularly appealing because it nests the null hypothesis of no private information ($F(p|X) = 1\{p \leq \Pr\{U|X\}\}$).³⁴

A key difficulty with using functions to approximate the distribution of P is that much of the mass of the distribution is near zero. Continuous probability distribution functions, such as the Beta distributions used in Hendren (2013b), require very high degrees for the shape parameters to acquire a good fit. Therefore, I approximate P as a sum of discrete point-mass distributions.³⁵ Formally, I assume

$$\tilde{F}(p|q) = w1\{p \leq q - a\} + (1 - w)\sum_i \xi_i 1\{p \leq \alpha_i\}$$

³³Specifically, the p.d.f./p.m.f. of Z given P is given by

$$f(Z|P, X) = \begin{cases} (1 - \lambda)\Phi\left(\frac{-P - \alpha(X)}{\sigma}\right) + \lambda\Phi\left(\frac{\kappa - P - \alpha(X)}{\sigma}\right) & \text{if } Z = 0 \\ \lambda\left(\Phi\left(\frac{1 - \kappa - P - \alpha(X)}{\sigma}\right) - \Phi\left(\frac{\kappa - P - \alpha(X)}{\sigma}\right)\right) & \text{if } Z = 0.5 \\ (1 - \lambda)\Phi\left(\frac{1 - P - \alpha(X)}{\sigma}\right) + \lambda\left(1 - \Phi\left(\frac{1 - \kappa - P - \alpha(X)}{\sigma}\right)\right) & \text{if } Z = 1 \\ \frac{1}{\sigma}\phi\left(\frac{Z - P - \alpha(X)}{\sigma}\right) & \text{if } o.w. \end{cases}$$

where ϕ denotes the standard normal p.d.f. and Φ the standard normal c.d.f. I estimate four elicitation error parameters: $(\sigma, \lambda, \kappa, \alpha(X))$. σ captures the dispersion in the elicitation error, λ is the fraction of focal point respondents, κ is the focal point window. I allow the elicitation bias term, $\alpha(X)$, to vary with the observable variables, X . This allows elicitation to be biased, but maintains the assumption that true beliefs are unbiased.

³⁴Moreover, it allows the statistical model to easily impose unbiased beliefs, so that $\Pr\{U|X\} = E[P|X]$ for all X .

³⁵This has the advantage that it does not require integrating over high degree of curvature in the likelihood function. In practice, it will potentially under-state the true variance in P in finite sample estimation. As a result, it will tend to produce lower values for $T(p)$ than would be implied by continuous probability distributions for P since the discrete approximation allows all individuals at a particular point mass to be able to perfectly pool together when attempting to cover the pooled cost of worse risks.

where α_i are a set of point masses in $[0, 1]$ and ξ_i is the mass on each point mass. I estimate these point mass parameters using maximum likelihood estimation. For the baseline results, I use 3 mass points, which generally provides a decent fit for the data. I then compute the pooled price ratio at each mass point and report the minimum across all values aside from the largest mass point. Mechanically, this has a value of $T(p) = 1$ – as noted in Hendren (2013b), estimation of the minimum $T(p)$ across the full support of the type distribution is not feasible because of an extremal quantile estimation problem. To keep the estimates “in-sample”, I report values for the mean value of $q = \Pr\{U\} = 0.031$; but estimates at other values of q are similarly large.

Results Table III reports the results for the same specifications and samples used in Table II for the lower bound estimates. I estimate a value of $\inf T(p) - 1$ of 3.36 in the baseline specification. This suggests that unless people are willing to pay a 336% markup in order to obtain unemployment insurance, the results are consistent with the absence of a private market. Including health controls reduces this markup slightly to 323%, but the frictions imposed by private information are quite large.³⁶ As shown in Appendix Table II and consistent with Figure II (Panel B), I estimate that there is a non-trivial fraction of the sample that has a very high chance of losing their job. The presence of this upper tail of high risks makes it incredibly difficult to profitably sell unemployment insurance.

The results are also quite robust across subsamples, as illustrated in Columns (4)-(9) of Table III. Consistent with the findings in the lower bound analysis, I find larger barriers to trade imposed by private information for those with longer tenure backgrounds (and hence lower unemployment probabilities on average), with values of $\inf T(p) - 1$ of 473.6%. The results are similar across age groups (3.325 for ages at or below 55 and 3.442 for ages above 55); and they are slightly higher for below-median wage earners (4.217) than above-median wage earners (3.223). Overall, the results suggest private information imposes a significant barrier to the existence of a private unemployment insurance market.

For comparison, Hendren (2013b) uses the same empirical strategy to study whether private information prevents those with pre-existing conditions from being able to purchase insurance in three market settings: Long-Term Care insurance, Life insurance, and Disability insurance. In those settings, the estimated markups are all below 100%: 42% for Life, 66% for Disability,

³⁶Consistent with the idea that the predicted values are lower bounds on the true distribution of beliefs, the estimated minimum pooled price ratios are all well in excess of the estimated lower bounds, $E[T_Z(P_Z)]$.

and 83% for Long-Term Care.³⁷ The size of the barrier to trade imposed by private information about future unemployment risk appears to be quite substantial.

4.3 Comparison to Willingness to Pay

How much of a markup would individuals be willing to pay for unemployment insurance after learning their type, θ ? Recall the individuals' willingness to pay is given by their marginal rate of substitution, $\frac{u'(c_u(\theta))}{v'(c_e(\theta))}$, where $c_u(\theta)$ and $c_e(\theta)$ are the consumption of a type θ in the event he or she is unemployed or employed. Following Baily (1976) and making the assumption that utility over consumption is state independent, $v = u$, one can use a Taylor expansion for u' around the consumption when employed, $u'(c) \approx u'(c_e(\theta)) + u''(c_e(\theta))(c - c_e(\theta))$, to yield the approximation:

$$\frac{u'(c_u(\theta))}{v'(c_e(\theta))} \approx 1 + \sigma \frac{\Delta c}{c}(\theta) \quad (11)$$

where $\frac{\Delta c}{c} = \frac{c_e(\theta) - c_u(\theta)}{c_e(\theta)}$ is the causal effect of the event of unemployment on type θ 's percentage difference in consumption and σ is the coefficient of relative risk aversion, $\sigma = \frac{c_e(\theta)u''(c_e(\theta))}{u'(c_e(\theta))}$. Gruber (1997) estimates this consumption drop between 5 and 10%. In Part II, I replicate these consumption patterns in the Panel Study of Income Dynamics using data up to 1997 and find a consumption drop of roughly 8% (see Figure VII). Assuming a coefficient of relative risk aversion of 2, it suggests individuals should be willing to pay roughly a 15-20% markup in order to buy a private unemployment insurance contract.³⁸

In short, the patterns of willingness to pay are consistent with the absence of the existence of a private market for UI. Although such willingness to pay would increase in the absence

³⁷Appendix Figure II illustrates this comparison.

³⁸Formally, this suggests individuals are not willing to pay to overcome the hurdles imposed by private information for additional insurance beyond what is currently provided in the status quo world by the government, their firms, friends and family, and other sources of formal and informal insurance. Indeed, the distribution of beliefs, P , in the status quo world are precisely what is desired for measuring whether a private market for additional unemployment insurance would arise. But, it is also natural to ask whether a private market would arise if the government were to lower the amount of UI it provides.

To address this, Gruber (1997) also explores how this consumption drop varies with the level of government unemployment benefits. Extrapolating to a world where the government provides no unemployment benefits, he shows the consumption drop would be roughly 25% (Table I, p196). This would imply individuals would be willing to pay a 75% markup for insurance if they had a coefficient of relative risk aversion of $\sigma = 3$. This value continues to be of the order of magnitude of the estimated lower bounds for $E[T(P)]$ and falls well below the estimated 300%+ markups for the point estimates for $\inf T(p)$ in Section 4.2. In principle, changing the amount of government benefits could change the markups imposed by private information, $T(p)$; however, the underlying fact that there appears to be a small fraction of people in every observable subgroup of the population that knows they are likely to lose their job would likely not be heavily affected; if anything, one might expect lower mean rates of unemployment entry which, as shown in Figure III, Panel F, would lead to higher markups that individuals would have to be willing to pay to cover the pooled costs of worse risks.

of government provision, the frictions imposed by private information are likely to continue to prevent the existence of a private market. The patterns are consistent with private information being a micro-foundation for the absence of a private unemployment insurance market and potentially present a rationale for government intervention.

Part II: Implications for Optimal Government Intervention

5 A Modified Baily-Chetty Condition

If private information prevents the existence of a private UI market, then no one is willing to pay the pooled cost of worse risks in order to obtain additional insurance. Additional UI benefits would not deliver a Pareto improvement – some types θ (e.g. the “good risks”) would be worse off, whereas other types (e.g. the “bad risks”) would be better off.³⁹

However, the endowment is not the only constrained-efficient allocation. A government can force the good risks to pay for insurance and accept utility levels below their endowment with no insurance. Traditional analyses of optimal social insurance solves for the optimal utilitarian policy – the level of benefits that maximizes the average level of utility across types, θ . This utilitarian metric can also be motivated from an ex-ante perspective of what level of UI benefits individuals would prefer prior to learning their type θ .

Before considering the optimal UI benefits in the present context, it is useful to begin with the canonical welfare analysis of UI. In a model without private information, Baily (1976) shows that the optimal level of UI benefits solves the optimality condition:

$$\frac{u' - v'}{v'} \approx \sigma \frac{\Delta c}{c} = FE \quad (12)$$

The optimal UI benefit level trades off the consumption-smoothing benefits of UI against its costs resulting from moral hazard.⁴⁰ The RHS of equation (12), FE , is the fiscal externality imposed by the behavioral response of individuals to the additional government UI benefits.⁴¹ The LHS of equation (12) is the difference in marginal utilities between being employed and unemployed, $\frac{u' - v'}{v'}$. It captures the individuals’ willingness to pay out of income in the employed state (at

³⁹Although formally the no trade condition only considers single contracts, Appendix A.2 illustrates that the no trade condition also rules out menus of contracts so that there cannot be Pareto improvements from menus of insurance contracts either.

⁴⁰In the absence of a micro-foundation for market non-existence, equation (12) characterizes not only the optimal insurance provided by the government, but also the optimal insurance provided by a competitive market. If $\frac{u' - v'}{v'} > FE$, it would suggest private firms should be able to profitably provide additional insurance. Of course, the presence of private information imposes a wedge, $T(p)$, as measured in Part 1, between what the government and private markets can provide.

⁴¹It is often written as $FE = \frac{\epsilon}{1-p}$ where ϵ is the duration elasticity of unemployment and p is the probability of employment; more generally, the fiscal externality is simply the causal impact of the behavioral response to additional benefits, b , financed by taxes, τ , on the government’s budget constraint (Chetty (2006); Hendren (2013a)). This in principle includes impacts from extensive margin entry into unemployment (Feldstein (1978); Topel (1983)), improved wages from increased job match quality (Schmeider et al. (2013); Nekoei and Weber (2015)), or fiscal impacts from changes in precautionary savings behavior (Engen and Gruber (2001)) or spousal labor supply (Cullen and Gruber (2000)).

utility cost v') to move resources to the unemployed state, valued at u'). Under the assumption that utility is state-independent, these differences can be approximated as $\sigma \frac{\Delta c}{c}$, where σ is the coefficient of relative risk aversion, $\sigma = -\frac{v''c}{v'} = \frac{u''c}{u'}$, and $\frac{\Delta c}{c}$ is the proportional “drop in consumption caused by a lay-off” (Baily (1976), p. 389). Equation (12) has led to a large literature focused on estimating these consumption drops upon unemployment and how they vary with the level of government benefits (e.g. Gruber (1997)). However, valuing the benefits of UI in terms of the causal effect of unemployment on consumption – or more generally, the causal effect of unemployment on marginal utilities – does not readily generalize to the case when individuals have knowledge about their future likelihood of job loss.

If individuals learn about future job loss, they can mitigate the causal impact of the event on their marginal utilities. Below, I show people respond to learning about *potential* job loss by cutting consumption and increasing spousal labor supply. To see how this can bias the calculation in equation (12), note that in the limit someone could take costly actions to fully smooth the onset of the event of unemployment if it were to occur, so that $u' = v'$. Hence, an econometrician looking at the causal impact of unemployment on marginal utilities would observe no consumption smoothing value of UI. But, from the individuals’ perspective, they would have wanted insurance against the event of learning they were going to lose their job.

Although the presence of knowledge about future events renders the causal effect of unemployment on marginal utilities a potentially biased measure of the value of social insurance, it does open up new directions for valuing social insurance. One can instead analyze the impact of information about unemployment on marginal utilities, as opposed to the impact of unemployment itself.

5.1 Setup

To begin, consider the optimal government benefit level in the model in Section 2. Let b denote the level of government unemployment insurance benefits, financed with taxes τ . The government seeks to maximize a utilitarian welfare function,

$$Q(\tau, b) = E[U(\tau, b; \theta)]$$

subject to the budget constraint

$$E[1 - p(\theta)]\tau - E[p(\theta)]b + E[N(a(\theta))] = 0$$

where $E[p(\theta)]b$ are the unemployment insurance payments, $E[1 + p(\theta)]\tau$ are the taxes collected from the employed to pay for the unemployment benefits, and $E[N(a(\theta))]$ is a placeholder that captures the net government budget from all other aspects of the individual's behavior (captured in $a(\theta)$).⁴² In contrast to markets, the government need not respect any participation constraint: it can force everyone to pay premiums, τ , so that the budget constraint involves the entire population, as opposed to the adversely selected subset, $\bar{\Theta}\left(\frac{d\tau}{db}\right)$.

It is straightforward to show that the level of b and τ that maximizes utilitarian welfare solves the modified Baily-Chetty condition:

$$W = \frac{E\left[\frac{p(\theta)}{E[p]}u'(c_u(\theta))\right] - E\left[\frac{(1-p(\theta))}{E[1-p]}v'(c_e(\theta))\right]}{E\left[\frac{(1-p(\theta))}{E[1-p]}v'(c_e(\theta))\right]} = FE \quad (13)$$

where FE is the fiscal externality associated with the policy.⁴³ The intuition for the difference

⁴²I include this term to illustrate that the FE component of the Baily formula remains in this more general setup. For example, if $a(\theta)$ includes spousal labor supply, N would include the net taxable income implications of this labor supply. If individuals can make choices that affect their future wages, N would include the net taxable income implications of those decisions.

⁴³To see this, note that the optimal allocation solves the first order condition:

$$\frac{\partial V}{\partial b} + \frac{\partial V}{\partial \tau} \frac{d\tau}{db} = 0$$

where

$$\frac{d\tau}{db} = \frac{E[p(\theta)]}{1 - E[p(\theta)]} + \frac{d}{db} \left[\tau \frac{E[p(\theta)]}{1 - E[p(\theta)]} + T(a(\theta)) \right]$$

is the increased premium required to cover the cost of additional benefits, which includes the impact of the behavioral response, $\frac{d}{db} \left[\tau \frac{E[p]}{1 - E[p]} + T(a(\theta)) \right]$. Note this includes the response from additional unemployment entry (e.g. $\frac{dE[p]}{db}$) and through any other behavioral response through changes in the choice of $a(\theta)$. Also, note these responses are “policy responses” as defined in Hendren (2013a) – they are the behavioral response to a simultaneous increase in b and τ in a manner for which the government's budget breaks even.

Now, one can recover the partial derivatives using the envelope theorem:

$$\frac{\partial V}{\partial b} = E[p(\theta)u'(c_u(\theta))]$$

$$\frac{\partial V}{\partial \tau} = -E[(1 - p(\theta))v'(c_e(\theta))]$$

So, the optimality condition becomes:

$$\frac{E\left[\frac{p(\theta)}{E[p(\theta)]}u'(c_u(\theta))\right]}{E\left[\frac{(1-p(\theta))}{E[1-p(\theta)]}v'(c_e(\theta))\right]} = 1 + FE$$

where

$$FE = \frac{\frac{d}{db} \left[\tau \frac{E[p(\theta)]}{1 - E[p(\theta)]} + N(a(\theta)) \right]}{\frac{E[p(\theta)]}{1 - E[p(\theta)]}}$$

If only p is the margin of adjustment, then

$$FE = \tau \frac{\frac{d}{db} \left[\frac{E[p(\theta)]}{1 - E[p(\theta)]} \right]}{\frac{E[p(\theta)]}{1 - E[p(\theta)]}} = \frac{\epsilon_{p,b}}{1 - E[p(\theta)]}$$

between equations (13) and (12) is straightforward. The envelope theorem implies individuals value additional benefits using their marginal utilities. The marginal utility of additional benefits to a type with probability $p(\theta)$ of experiencing unemployment is pu' . The cost to the government of providing an additional dollar of benefits is proportional to the average probability of unemployment in the population, $E[p]$. If individuals are identical in their probabilities of experiencing unemployment (e.g. no one has unique knowledge about the event), then $p = E[p]$, and the formula reduces to the canonical formula in equation (12) with the average utilities $E[u']$ and $E[v']$ in place of u' and v' . Under state independence, $u' = v'$, one can then approximate W as $\sigma \frac{\Delta c}{c} = FE$, where σ is the coefficient of relative risk aversion and $\frac{\Delta c}{c}$ is the average causal effect (i.e. within- θ difference) of unemployment on consumption,

$$\frac{\Delta c}{c} = \frac{E[c_e(\theta) - c_u(\theta)]}{E[c_e(\theta)]}$$

But more generally, if individuals respond to their knowledge of future unemployment and take actions to mitigate its impact on their marginal utilities, this will introduce a covariance between p and the marginal utility of consumption, u' . In this case, the causal impact of unemployment on marginal utilities (i.e. $\frac{u'(c_u(\theta))}{v'(c_e(\theta))}$) is no longer sufficient for measuring W .⁴⁴

where $\epsilon_{p,b}$ is the elasticity of the unemployment probability with respect to the benefit level. More generally one would need to incorporate the fiscal externality associated with the responses from a (e.g. wages).

⁴⁴Although the causal effect is no longer sufficient when individuals have ex-ante information, Chetty (2006) considers a generalized model where individuals may know about their future unemployment prospects. In this case, he shows that one can approximate equation (13) using the ex-ante expected difference in consumption

$$\frac{E\left[\frac{p(\theta)}{E[p]}u'(c_u(\theta))\right]}{E\left[\frac{(1-p(\theta))}{E[1-p]}v'(c_e(\theta))\right]} \approx \sigma \frac{\Delta c^{lifetime}}{c}$$

where

$$\frac{\Delta c^{lifetime}}{c} \approx \frac{E\left[\frac{U c_u(\theta)}{\Pr\{U=1\}}\right] - E\left[\frac{1-U}{\Pr\{U=0\}}c_e(\theta)\right]}{E\left[\frac{1-U}{\Pr\{U=0\}}c_e(\theta)\right]} = \frac{E\left[\frac{p(\theta)}{E[p]}c_u(\theta)\right] - E\left[\frac{1-p(\theta)}{1-E[p]}c_e(\theta)\right]}{E\left[\frac{1-p(\theta)}{1-E[p]}c_e(\theta)\right]}$$

Note this is not the causal impact of the event of unemployment on consumption. If individuals who experience unemployment tend to have lower consumption in both states of the world then this number would not be zero even if $c_e(\theta) = c_u(\theta)$. If one had a sample of ex-ante identical individuals, $\frac{\Delta c^{lifetime}}{c}$ would be identified from the average (cross-sectional) difference in consumption in the employed and unemployed. Note this would capture mitigation responses because those who learn they have a high likelihood of becoming unemployed may have lower consumption in both the employed and unemployed state. But, one cannot look within the set of unemployed and calculate their consumption that they would have experienced had they not been employed to calculate $\frac{\Delta c^{lifetime}}{c}$. One could potentially use data in the cross-section to compare employed and unemployed, but this would mix uncertainty with heterogeneity.

5.2 Ex-Ante and Ex-Post Behavior

If individuals learn about future unemployment before the event occurs, it should effect behavior at the point when they learn (ex-ante relative to the job loss). Moreover, it should effect outcomes after the uncertainty over job loss is realized, $c_u(p)$ and $c_e(p)$. Theory suggests that these ex-ante and ex-post behaviors are linked through an Euler equation. Let $c_{pre}(p)$ denote the consumption of an individual at the time when learning θ .⁴⁵ For simplicity, assume the mapping from θ to p is 1-1 so that the heterogeneity is fully characterized by p . The individual's choice of how much to consume today relative to saving to consume after the event of unemployment is realized will satisfy an Euler equation

$$v'(c_{pre}(p)) = pu'(c_u(p)) + (1-p)v'(c_e(p)) \quad (14)$$

Because the marginal utility of consumption today will vary with their beliefs about losing their job in the future, individuals may make choices that affect their consumption in the future, $c_u(p)$ and $c_e(p)$. These choices could include spousal labor supply, savings, informal insurance mechanisms, job search, etc.

The Euler equation suggests there are two methods for understanding the degree to which individuals' behavior is affected by future beliefs, p . First, one can ask whether the ex-ante information affects ex-post allocations, $c_e(p)$ and $c_u(p)$. If so, then the canonical welfare formulas do not capture the ex-ante value of social insurance. Second, one can measure the impact of learning about future unemployment on ex-ante choices prior to the event. I consider these in turn and provide new measures of the social value of unemployment insurance.

6 Does information affect ex-post allocations, c_u and c_e ?

6.1 Data

I begin by exploring whether knowledge about future unemployment impacts consumption after the event of unemployment is realized, $c_u(p)$ and $c_e(p)$. To do so, I rely on the consumption mail survey component of the HRS, which is mailed to roughly 10% of respondents. It provides information on a range of consumption variables that are aggregated in a cross-year file constructed by RAND. It is administered 1 year after the core survey and asks about consumption

⁴⁵Note c_{pre} is captured in the model in Section 2 by thinking of it as an element of a .

expenditure in the previous 12 months. Hence, it provides a measure of consumption in the time period corresponding to the unemployment measure, U .

Table I, Panel 3 presents the summary statistics for the consumption sample. There are 2,798 observations from 862 households. The consumption module is asked of the entire household. To account for differences in household size, I present results for both aggregate household consumption and per capita consumption, which is household consumption divided by the total number of household members. All standard errors are clustered at the household level.

6.2 Results

Figure IV plots the relationship between group indicators of the subjective probability elicitation and log per capita consumption expenditure (Panel A) and log consumption expenditure (Panel B). The regression includes controls for census region, year, age, age squared, gender, marital status, the log wage, and – most importantly – an indicator for the future realization of unemployment. As shown in the figure, there is an decreasing pattern over the range $Z > 0$: individuals with higher subjective probability elicitation have lower consumption. Consistent with the measurement error model in Section 4.2 that suggests most of the reports of $Z = 0$ reflect a point bias that would have otherwise had higher values of Z , we obtain a lower coefficient at $Z = 0$.

Motivated by the non-parametric pattern in Figure IV, Table IV reports the regression coefficient on Z , combined with a dummy indicator for $Z = 0$. These variables are interacted with an indicator of subsequent unemployment U . Column (1) reports the negative coefficient of -0.16 (s.e. 0.0781) for the per capita consumption specification for those who do not experience unemployment. Those who believe they are more likely to become unemployed have lower consumption even if they do not become unemployed. This pattern is precisely what can lead to the canonical Baily formula under-stating the value of social insurance. The coefficient on the interaction with unemployment is negative, -0.137 (s.e. 0.268), but not statistically significant. This should not be too surprising given the fact that roughly 3% of households actually experience unemployment. The negative coefficient on $1\{Z = 0\}$ of -0.0893 (s.e. 0.0334) is consistent with the pattern in Figure IV in which the consumption expenditure values at $Z = 0$ fall below the pattern generated by the positive elicitation.

Column (2) reports the results using household consumption instead of household consump-

tion per capita. This yields a coefficient of -0.110 (s.e. 0.0596) on the elicitation for those who do not experience unemployment. Here, the coefficient on the unemployment interaction with the elicitation is statistically significant, -0.421 (s.e. 0.207), but is arguably too large for credibility and has a very wide standard error. Column (3) restricts the sample to those who have positive elicitations and illustrates that the results are quite similar to the baseline specification in Column (1). Column (4) restricts the sample to those who do not experience unemployment; here again, the coefficients are similar to the baseline specification. Column (5) considers non-durable consumption instead of total consumption expenditure, and finds a negative coefficient of -0.162 (s.e. 0.0789) that is again similar to the baseline specification.

Column (6) drops the control variables in the analysis. Here, we end up with a larger coefficient of -0.345 (s.e. 0.0798) from the analysis. A key concern with this specification is that the variation in beliefs captures heterogeneity in people (e.g. low versus high wage workers) as opposed to learning about the event. I return to the distinction between selection and information realization below.

Finally, column (7) illustrates the fragility of the results to the inclusion of the indicator for $Z = 0$. As shown in Figure IV, the negative relationship is quite nonlinear. While this pattern is consistent with focal point bias so that many of those responding $Z = 0$ are drawn from a population who otherwise would have said a much larger value of Z , dropping these controls renders the negative slope insignificant at -0.04 (s.e. 0.0659).

Selection versus the effect of information realization The cross-sectional relationship between the ex-ante subjective elicitations and consumption could reflect either the impact of learning about future unemployment on consumption, or be the result of a general correlation across the income distribution between job stability and income levels. Although the regressions control for the individual's wages, there of course could be measurement error in the survey, or it could be that many years of lagged wages are relevant.⁴⁶

To disentangle whether the patterns in Figure IV and Table IV reflect an impact of information revelation about future unemployment or a cross-sectional selection pattern, Figure V,

⁴⁶If the pattern reflects selection between high and low income individuals, the covariance calculations for the optimal degree of unemployment insurance would be valid as measuring the benefits from additional UI, but one would want to take into account the impact of UI on the effective total amount of redistribution in the economy and include the associated fiscal externalities akin to the redistributive costs associated with the progressive income tax schedule (Kaplow (2008); Hendren (2014)).

Panel A, presents the coefficients on the elicitation, Z , using leads and lags of log household consumption expenditure per capita. I include controls for unemployment status and an indicator for a focal point response of $Z = 0$. For simplicity, I summarize the negative relationship between Z and log consumption expenditure by pooling across unemployment status and do not interact Z with U .

Figure V, Panel A, reveals that higher values of subjective probability elicitation do not correspond to lower values of consumption when measuring consumption in the years prior to the elicitation (conditional on the controls for census region, year, age, age squared, gender, marital status, and the log wage). Rather, the onset of the realization of information about a greater likelihood potential unemployment leads to lower consumption in the years subsequent to the information revelation. This is consistent with the idea that the pattern in Figure IV is largely capturing the impact of information shocks on consumption, as opposed to a persistent heterogeneity in consumption across the population who report high versus low elicitation, Z .

Panel B replicates the analysis on the subsample with positive elicitation only ($Z > 0$), corresponding to column (3) in Table IV. Panels C and D replicate the analysis using household consumption instead of per capita consumption. Across all specifications, we find the pattern that consumption appears to drop at the point of learning about future likely unemployment, even conditional on whether or not that unemployment actually occurs.

6.3 Bounds on Welfare

To quantify the magnitude of the impact of beliefs on consumption, one can relate the coefficients to the implications for the willingness to pay for social insurance. A key difficulty in doing so is the measurement of the marginal utility of consumption when unemployed, $u'(c_u)$. There is not much data for whom $U = 1$, especially given the small size of the consumption survey it is difficult to get reliable estimates on this subsample. Moreover, even armed with good estimates one may worry that utility is state-dependent, so that $u' \neq v'$.

To circumvent these issues and arrive at a welfare implication for the slope in Table IV, one can use the Euler equation (14) to replace $c_u(p)$ with $c_{pre}(p)$. In doing so, Appendix C.1 shows that a Taylor approximation yields:

$$W \approx \frac{1}{E[p]} \sigma \frac{\Delta c^{NotU}}{c} + \sigma cov \left(\frac{-p}{E[p]}, \frac{c_e(p)}{E[c_e(p)]} \right)$$

where $\frac{\Delta c^{NotU}}{c} = \frac{E[c_e(p)] - E[c_{pre}(p)]}{E[c_{pre}(p)]}$ is the average impact on consumption in the event of not

losing ones' job. In general, one would expect that not losing ones' job leads to an increase in consumption, so that $\frac{\Delta c^{NotU}}{c} \geq 0$. Hence, one can arrive at a lower bound on W ,

$$W = \sigma \text{cov} \left(\frac{-p}{E[p]}, \frac{c_e(p)}{E[c_e(p)]} \right) \approx \sigma \frac{\text{var}(Z)}{\Pr\{L\}} \frac{\text{cov}(-Z, \log(c_e))}{\text{var}(Z)}$$

where $\frac{\text{cov}(-Z, \log(c_e))}{\text{var}(Z)}$ is the regression coefficient in on Z in Table IV. This formula suggests taking the regression coefficient (multiplied by -1) and multiplying by variance of the elicitation, Z , to arrive at the covariance of the elicitation and log consumption. Then, under the unbiased beliefs assumption, $\Pr\{L\} = E[p]$, yielding an estimate of W .

Panel 2 of Table IV presents these lower bound results for W . For the baseline specification and an assumption of $\sigma = 2$, I estimate individuals are willing to pay at least a 49% markup for unemployment insurance. Across the range of specifications, the markups range from 34% using household consumption instead of household consumption per capita, to 50% if one uses non-durable consumption per capita instead of durable consumption per capita. Dropping the controls increases the markup to 107%. Intuitively, such a policy would entail redistribution across high versus low wage populations, and hence would likely involve additional fiscal externalities that, in many cases would offset the value of the redistribution (Kaplow (2008)). Finally, Column (7) presents another reminder that the results are sensitive to the inclusion of the controls for the focal point response of $Z = 0$. Indeed, a drawback to the analysis of this section is that it relies on a fairly small sample and generates fairly large standard errors relative to the effect sizes. While the results suggest ex-post allocations vary with ex-ante beliefs, the next section turns to ex-ante responses to information that allow one to utilize larger sample sizes.

7 Ex-Ante Responses

Behavior should respond to information about future events at the time the individuals learn that information. Here, I document two such responses: spouses are more likely to enter the labor force and households lower their consumption.

While the traditional welfare metric W requires computation of ex-post behavioral responses, there is a natural analogous welfare metric that relies solely on ex-ante behavior at the time of learning about future unemployment. Consider the subset of people who learn their future

unemployment status with certainty, $p \in \{0, 1\}$. From the Euler equation (14),

$$v'(c_{pre}(1)) = u'(c_u(1))$$

$$v'(c_{pre}(0)) = v'(c_e(0))$$

so that

$$W^{ex-ante} = \frac{v'(c_{pre}(1)) - v'(c_{pre}(0))}{v'(c_{pre}(0))} \quad (15)$$

$W^{ex-ante}$ evaluates the willingness to pay across states of the world on the subset of the population that learns with certainty that they will or will not lose their job. In practice, it will be easiest to estimate $W^{ex-ante}$ using interior values of p as well, using the approximation

$$W^{ex-ante} \approx \frac{d \log(v'(c_{pre}(p)))}{dp} \approx \sigma_v \frac{d \log(c_{pre})}{dp} \quad (16)$$

where $\sigma_v = -\frac{v''c}{v'}$ is the coefficient of relative risk aversion (evaluated at $c_{pre}(0)$). $W^{ex-ante}$ measures the percentage change in ex-ante marginal utilities from a percentile increase in the likelihood of unemployment. It measures the value of moving resources from someone that has learned he or she is less likely to lose his or her job to someone who has learned he or she is more likely to lose his or her job. In this sense, the welfare metric is identified from the subset of the population that has more knowledge about their future unemployment. In general, one might expect individuals who learn more about the future to be able to react and make better decisions than those for whom unemployment comes as a surprise. This would lead one to expect $W^{ex-ante} \leq W$. In this sense, $W^{ex-ante}$ provides a lower bound on the welfare value of additional unemployment insurance. It has the key benefit, however, that it does not require assumptions about state dependence (v versus u).

This section presents two approaches to identifying ex-ante behavioral responses and estimating $W^{ex-ante}$. The key hurdle is that the HRS does not provide consumption measures concurrently with the elicitation measures, so direct computation of $c_{pre}(p)$ is not feasible. I follow two approaches. First, I estimate the impact on spousal labor supply and estimate the value of social insurance using the intra-temporal tradeoff between labor and leisure and an assumption about the elasticity of spousal labor supply. Second, I estimate equation (16) using a two-sample IV approach that uses the time evolution of knowledge about future unemployment as an instrument for consumption in the PSID (the reduced form) and for beliefs in the HRS (the first stage).

7.1 Impact on Spousal Labor Supply

When individuals learn they may lose their job, their marginal utility of income should increase. This should not only decrease consumption activities, but also it may increase activities that generate income, such as spousal labor supply. Here, I consider the impact on spousal labor entry into the labor market. Suppose a spouse can enter the labor market to earn $\$y$ at a fixed cost of η , where η is distributed heterogeneously in the population.⁴⁷ At an optimum, all individuals where the marginal utility of that $\$y$ exceeds η will choose to work. Appendix C.2 shows that

$$W^{ex-ante} \approx \frac{d\phi}{dp} \frac{1}{\epsilon^{semi}} \quad (17)$$

where $\frac{d\phi}{dp}$ is the percentage point increase in labor force participation resulting from a 1pp increase in p , and ϵ^{semi} is the semi-elasticity of spousal labor supply, equal to the percentage point increase in labor force participation that arises from a percentage increase in earnings. By taking the ratio of the impact of learning about unemployment relative to the impact of an increase in wages, equation (17) translates the impact on the labor force participation rate into units of the marginal utility of income.

Results Using the HRS sample of all households with spouses, I define labor market entry by the spouse as an indicator for the spouse working for pay in the current wave of the survey and not working for pay in the previous wave of the survey (2 years prior). I restrict the sample to individuals who are married in both the current and previous wave. Figure VI plots the coefficients on the subjective probability elicitation controlling for census region, year, age, age squared, gender, marital status, the log wage, and an indicator for the future realization of unemployment.

The figure suggests spouses of those with higher elicitation are more likely to enter the labor force. The pattern is slightly different than for consumption, with labor market entry more likely if the subjective elicitation is quite high. The spouses of individuals with $Z > 50$ as opposed to $Z = 0$ are 2 percentage points more likely to enter the labor force. On the one hand, this is a small effect: it suggests roughly 1 in 50 extra spouses are induced into the labor market when the spouse reported an elicitation above 50%, $Z > 50$. On the other hand, relative to the base entry rate of these spouses of 3.9%, it is quite large. For values $Z < 50$, the response is

⁴⁷Spousal labor supply is incorporated into the general model in Section 2 by thinking of it as an element of a .

more muted. In contrast, the consumption response shown in Figure IV is largely continuous in Z for positive values of Z below 50%, $Z < 50$. This is suggestive of a model in which labor market entry has higher fixed costs than small continuous modifications to consumption patterns, as would be implied by many labor market models. In response to smaller variations in information, individuals adjust consumption; once the chance of job loss crosses a sufficient threshold, spouses become more likely enter the labor market.

Table V provides a linear parameterization of the regression relationship. Column (1) of Table V presents this coefficient, and consistent with the approach for ex-post consumption in Table IV, also includes a control for $Z = 0$. This yields a slope of 0.0273 (s.e. 0.0112). In contrast to the ex-post consumption measures above, the results are robust to dropping the control for focal point elicitation of $Z = 0$, as shown in Column (2). Column (3) restricts the sample to those who do not end up losing their job in the 12 months after the survey, yielding similar results. Column (4) uses a specification that defines spousal work as an indicator for full-time employment, as opposed to any working for pay. This definition counts shifts from part time to full time work as labor market entry, as opposed to transitions to work for any pay. The pattern is largely similar, with a slope of 0.0286 (s.e. 0.0128).

One identification concern about the results could be that individuals who are more likely to lose their jobs also have spouses that perhaps have less labor force attachment and are more likely to come and go into the labor market. If true, it could generate a correlation between labor market entry and the elicitation purely because of a selection effect. To this aim, Column (5) considers a placebo test that uses the lagged measure of entry, which corresponds to the previous wave of the survey conducted 2 years prior. Here, the coefficient is 0.00792 (s.e. 0.0102) and is not statistically distinct from zero. Column (6) adds household fixed effects to the regression in Column (1) and Column (7) adds individual fixed effects to the specification in Column (1). The point estimates are quite stable, although noisy with the individual fixed effects. Overall, the pattern appears to be consistent with the hypothesis that households respond to learning about unemployment by increasing spousal labor supply.

In addition to impacts on entry, one may also expect to see fewer spouses leave the labor force in response to learning about future unemployment prospects for the other earner. However, one does not find much evidence of this pattern in the data. Column (8) defines labor market exit as an indicator for a spouse working for pay last wave and not working for pay in the current

period. In contrast to the idea that spouses would be less likely to choose to enter the labor market, the coefficient of 0.017 (s.e. 0.0116) is positive, although not statistically significant.

One possible explanation for why spouses are not less likely to stop working could be that it's not their own choice; if a husband is likely to lose his job, the same set of circumstances may also affect the ability of the wife to stay in her job. To this aim, Column (9) shows that the elicitation is positively related to spousal unemployment in the subsequent year. Spouses of those who learn they may lose their job may wish to keep their job, but may not always have that choice. In this case, the estimates for the impact of learning about future job loss on spousal labor supply under-state the response that would occur if the opportunity set available to the spouse were held fixed.

Welfare Panel 2 of Table V presents the estimated values of $W^{ex-ante}$ using equation (17). To do so, I not only divide by the semi-elasticity of labor supply (here assumed to be 0.5), but also correct for the fact that the regressions estimate $\frac{d\phi}{dZ}$ as opposed to $\frac{d\phi}{dp}$. Measurement error in Z induces a potential attenuation bias. To do so, I scale the estimates by $\frac{var(Z|X)}{var(P|X)}$, where X are the controls in the regressions of labor force participation on Z .⁴⁸

The results suggest that individuals would be willing to pay a 60% markup, $W^{ex-ante} \approx 60\%$, for insurance against the event of learning they are going to lose their job in the baseline specification. The other specifications generate similar measures, which is not surprising given the stability of the regression coefficient in Panel 1. Moreover, these estimates are lower bounds to the extent to which spousal labor supply is constrained from correlated shocks and to the extent to which $W^{ex-ante}$ is less than the value of insurance on the set that includes those that do not learn about their future unemployment, W .

For comparison, the results from Chetty (2008) that calculate the impact of additional benefits on unemployment duration suggests a value of $FE \approx 0.5$.⁴⁹ Assuming this size of a fiscal externality from additional benefits, the size of the spousal labor supply response suggests additional UI would be preferred from an ex-ante perspective.

⁴⁸To do so, I construct $var(Z|X)$ as the square of the RMSE of a regression of Z on the control variables. I construct $var(P|X)$ as $var(P|X) \approx cov(Z, L|X)$, where the approximation would hold exactly if the measurement error in Z were classical. To construct $cov(Z, L|X)$, I first residualize L and Z on X and then calculate the covariance of the residuals, then adjust for the degrees of freedom introduced in the initial residualization.

⁴⁹Note this behavioral response is not in terms of the response of probability of unemployment; but the choice of unemployment duration can be thought of as incorporated into other behavioral responses, $a(\theta)$. Later work has found mixed evidence on the impact of UI on wages in various contexts (Schneider et al. (2013); Nekoei and Weber (2015)), which would also be ideally incorporated into these analyses.

Overall, the spousal labor supply response is quite large and indicative of a significant value of social insurance. Moreover, the overall pattern is also consistent with the finding of Gruber and Cullen (1996) that higher levels of social insurance reduce the response of spouses into the labor market in response to unemployment. The presence of greater social insurance reduces the degree to which learning about future unemployment increases the marginal utility of income, which reduces the incentives to enter the labor force. Relative to this literature, I find that a large fraction of these responses occur even before the onset of unemployment.⁵⁰

7.2 Consumption

In addition to spousal labor supply responses, one would also expect individuals who learn about future unemployment to cut back on their current consumption. As noted above, a key hurdle to identifying these effects is the absence of concurrent consumption measures in the HRS. Here, I develop a 2-sample IV approach to estimating the impact of information about future unemployment on current consumption and use this measure to estimate $W^{ex-ante}$. The approach allows one to utilize the information about beliefs in the HRS combined with the information on consumption that is available in other datasets – in this instance, I use the Panel Study of Income Dynamics (PSID). For identification, I utilize the variation in beliefs induced by the passage of time between 1 and 2 years prior to the onset of unemployment.

Consumption Patterns in PSID Many papers studying optimal unemployment insurance have used the PSID to measure the impact of unemployment on consumption. Here, I revisit this dataset with the different focus of studying the impact of future unemployment on earlier consumption patterns, where the time elapsation relative to the future unemployment spell provides an instrument for the evolution of knowledge about the unemployment event.

I utilize data on food consumption for years spanning 1971-1997. I restrict the sample to heads of household between the ages of 25 and 65 who have non-missing food expenditure and employment status variables. I define food expenditure as the sum of food expenditure in the home and out of the home, plus food stamps.⁵¹ Following Gruber (1997), I restrict the sample

⁵⁰As shown in Column (3) of Table V, the response occurs even on the subset of those who find it likely they will lose their job but who do not actually end up losing their job in the 12 months after the survey.

⁵¹To compute food stamp expenditure, I follow previous literature and use the response to the monthly food stamp amount multiplied by 12. Results for the impact on consumption in $t - 2$ relative to $t - 1$ are robust to alternative measures of food stamps, such as using the annual measures. However, the size of the consumption drop upon unemployment is larger when using the annual food stamp expenditure question instead of the monthly

to those with less than a threefold change in food expenditure relative to the previous year. I define an indicator for unemployment at the time of the survey that exclude temporary layoffs. Appendix Table III provides the summary statistics for the sample. In contrast to the less than 1,000 households in the HRS with consumption data, the PSID offers more than 10,000 household observations with food consumption data in the primary sample.

I begin by exploring the pattern of consumption around unemployment spells. Following previous literature, for each year I construct the change in log food expenditure relative to the previous year, $g_t = \log(c_t) - \log(c_{t-1})$. Figure VII illustrates how food expenditure growth in year $t + j$, g_{t+j} , relates to the onset of unemployment in year t for $j = -4, \dots, 4$. I regress g_{t+j} on an indicator for unemployment in year t , controlling for a cubic in age and year dummy indicators. Panel A reports the pattern for the entire sample. Panel B restricts the sample to those who are not unemployed in years $t - 1$ and $t - 2$.

As has been well-documented in previous work, there is a large consumption expenditure drop upon unemployment. The coefficients at $j = 0$ illustrate a roughly 6-8% drop. But, consistent with the hypothesis that individuals can partially forecast their future unemployment, the onset of unemployment in year t is associated with a 2.5% lower consumption growth between year $t - 2$ and $t - 1$, despite those individuals not being unemployed in either of those periods.

To further explore the robustness of this pattern and quantify the magnitude of the expenditure drop in the 1-2 years prior to unemployment, Table VI presents the results of a regression of the difference in log food expenditure in year $t - 2$ and year $t - 1$, $\log(c_{t-1}) - \log(c_{t-2})$, on an indicator for unemployment in current period. Column (1) includes controls for only age and year dummies, analogous to the specification in Figure VII, Panel A. This shows a -0.0336 (s.e. 0.0057) drop in expenditure in the year before unemployment occurs. Column (2) restricts the sample to those who are not unemployed in years $t - 1$ and $t - 2$, analogous to the specification in Panel B of Figure VII. This attenuates the food expenditure drop slightly to -2.5% (s.e. 0.00942). This is to be expected given the auto-correlation of 0.387 in unemployment status across years in the baseline sample.

An additional concern is that household size or needs change around the time of unemployment. Column (3) of Table VI adds controls for both the change in household size in years $t - 2$ versus $t - 1$ and also the change in expenditure needs, a variable available in the PSID response multiplied by 12.

that captures the total needs of the household based on its size and composition.⁵² This delivers a coefficient of -0.0249 (s.e. 0.0994), very similar to the -0.025 (s.e. 0.00942) coefficient in Column (2). Column (4) adds individual fixed effects to the specification in Column (2) and again delivers a coefficient of -0.0231 (s.e. 0.013), close to the -0.025 in Column (2). Column (5) restricts the sample to individuals over age 40 to more closely align with the HRS sample, which yields a coefficient of -0.0287 (s.e. 0.0151). Finally, Column (6) expands the sample to include outliers with more than threefold changes in food expenditure, yielding a coefficient of -0.0231 (s.e. 0.0121). Overall, the results suggest that future unemployment leads to a consumption drop of roughly 2.5% in the 1-2 years prior to the unemployment event.

Evolution of Beliefs To arrive at an estimate of $\frac{d \log(c_{pre})}{dp}$, the 2.5% consumption drop needs to be scaled by the amount by which information is revealed between 2 and 1 year prior to the onset of unemployment event. Let U_t denote an indicator for unemployment in year t . Let $P_{j,t}$ denote an indicator of the individuals beliefs at time $j \leq t$ about becoming unemployed in year t . The amount of information that is revealed by becoming unemployed in year $t - 1$ relative to $t - 2$ is given by:

$$\Delta^{\text{First Stage}} = \underbrace{E[P_{t-1,t}|U_t = 1] - E[P_{t-1,t}|U_t = 0]}_{\text{Knowledge in } t-1 \text{ about } t} - \underbrace{E[P_{t-2,t}|U_t = 1] - E[P_{t-2,t}|U_t = 0]}_{\text{Knowledge in } t-2 \text{ about } t}$$

The first component of $\Delta^{\text{First Stage}}$ can be obtained by simply regressing the elicitation, Z , on an indicator for unemployment in the subsequent 12 months, U . The first row of Appendix Table IV reports these results and shows this coefficient is 0.197 (s.e. 0.0123). Under the assumption that the measurement error in Z is uncorrelated with the subsequent event, U , this provides an unbiased measurement of $E[P_{t-1,t}|U_t = 1] - E[P_{t-1,t}|U_t = 0]$. It suggests that in the year prior to unemployment, the average difference in the beliefs, P , for those who subsequently do experience unemployment versus those that do not subsequently experience unemployment is roughly 20pp.

However, some of this difference in beliefs was already present in the 2 years prior. To subtract off the value of $E[P_{t-2,t}|U_t = 1] - E[P_{t-2,t}|U_t = 0]$, one would ideally have an elicitation about unemployment in the 12-24 months after the elicitation. However, we can get a lower

⁵²For some years, the PSID also has food need measures available. Controlling for these reduces the sample size but delivers similar results.

bound by using the elicitation about the future 12 month unemployment to predict unemployment in the 12-24 months after the survey. The second row of Appendix Table IV reports this coefficient as 0.0937 (s.e. 0.0113). Intuitively, people know more about their prospects for losing their job in the next 12 months than in the 12-24 months from now. The results suggest this difference-in-difference in beliefs between the unemployed and employed in years $t - 2$ and $t - 1$ is 0.1031 (s.e. 0.0159), as shown in the first row of Panel 2.

The next row scales the reduced form coefficients in Panel 1 by the first stage difference in beliefs of 0.1031. For the baseline specification in Column (2) using the sample that are employed in years $t - 2$ and $t - 1$, this yields a value of $\frac{d \log(c_{pre})}{dp} = 0.24$ (s.e. 0.09). This suggests learning one is 10% more likely to lose their job would cause a 2.4% drop in consumption. Scaling this estimate by a coefficient of relative risk aversion of $\sigma = 2$, it implies $W^{ex-ante} = 0.48$ (s.e. 0.18), which suggests individuals would be willing to pay a 48% markup for additional unemployment insurance. The remaining columns use the same first stage estimation to scale the estimates for the other specifications. These results generally fall around 50%. Compared to a value of FE of 0.5, they suggest the current level of unemployment benefits is about optimal provided $W \approx W^{ex-ante}$. If the value of UI to those that learn about their future unemployment is lower than to those for whom it comes as a complete surprise, then the point estimates would suggest $W > W^{ex-ante} \approx 0.5$, so that additional unemployment insurance benefits would be ex-ante preferred.

8 Threats to Identification

There are several threats to identification to keep in mind when interpreting the results. Throughout, one must make the assumption that variation in the elicitation represents measures of beliefs and is not correlated with changes in underlying preferences. For example, individuals could perhaps face time-varying preference shocks to the marginal utility of consumption and, in response to such a shock that lowers marginal utilities of consumption choose both (a) to put in less effort at work (and hence have a higher p) and (b) to consume less. While possible, the finding that spousal labor supply also increases in response to higher p suggests that indeed the marginal utilities of income are increasing, not decreasing, in response to p .

Second, the formula is derived in a static setting. Formally, it measures the individual's ex-ante willingness to pay (prior to learning θ) for an unemployment insurance contract that

pays \$1 in the event he or she lose her job in the subsequent 12 months. A conceptually distinct, but arguably more important question is whether the government should increase UI benefits in the steady state of a dynamic model. I leave the derivations of these adjustment factors to future work.

Finally, it is important to note that, as with the standard Baily formula, the ultimate welfare conclusions depend on the coefficient of relative risk aversion, σ , or the extensive margin labor supply elasticity, ϵ^{semi} . However, in contrast to the traditional approach looking at the impact of unemployment, here these coefficients are defined within the employed state of the world. In this sense, the approach developed here that exploits information revelation prior to the actual event is arguably less likely to suffer from bias arising from state-dependent marginal utilities.

9 Conclusion

This paper argues that private information prevents the existence of a robust private market for unemployment or job-loss insurance. Unless individuals are willing to pay extreme markups, the empirical results are consistent with the absence of a private market.

This micro-foundation motivates a modification to the formulas characterizing the utilitarian-optimal unemployment insurance benefit level. If individuals learn about unemployment before it actually occurs, they may respond to mitigate the causal effect of unemployment on their marginal utilities of income. I provide evidence that individuals respond by lowering consumption and increasing spousal labor force entry. While this renders traditional welfare analyses that focus on the causal effect of unemployment insufficient for welfare analysis, I illustrate how to use the response to information about future unemployment to estimate the ex-ante value of unemployment insurance.

The approaches can be applied to other settings, such as disability insurance, social security, and health insurance contexts. In particular, the 2-sample IV procedure developed in Section 7.2 shows one can conduct such welfare analysis without necessarily simultaneously observing consumption and beliefs. In this sense, I hope the analyses in this paper provides a path forward for motivating a micro-foundation for government intervention and the calculation of the optimal government intervention.

References

- Acemoglu, D. and R. Shimer (1999). Efficient unemployment insurance. *Journal of Political*. 7
- Acemoglu, D. and R. Shimer (2000). Productivity gains from unemployment insurance. *European Economic Review* 44(7), 1195–1224. 7
- Aguiar, M. and E. Hurst (2005). Consumption versus expenditure. *Journal of Political Economy* 113(5), 919–948. 1
- Akerlof, G. (1970). The market for lemons: Qualitative uncertainty and the market mechanism. *Quarterly Journal of Economics* 84(3), 488–500. 2.2
- Andrews, I. and C. Miller (2013). Optimal social insurance with heterogeneity. *MIT Working Paper*. 8
- Baily, M. (1976). Some aspects of optimal unemployment insurance. *Journal of Public Economics* 10, 379–402. 4, 1, 7, 4.3, 5, 5
- Barceló, C. and E. Villanueva (2010). The response of household wealth to the risk of losing the job: Evidence from differences in firing costs. 9
- Bloemen, H. G. and E. G. Stancanelli (2005). Financial wealth, consumption smoothing and income shocks arising from job loss. *Economica* 72(287), 431–452. 9
- Browning, M. and T. Crossley (2001). Unemployment insurance levels and consumption changes. *Journal of Public Economics* 80, 1–23. 1
- Cabral, M. (2013). Claim timing and ex post adverse selection. *Working Paper*. 31
- Carroll, C. D., K. E. Dynan, and S. D. Krane (2003). Unemployment risk and precautionary wealth: Evidence from households’ balance sheets. *Review of Economics and Statistics* 85(3), 586–604. 9
- Carroll, C. D. and A. A. Samwick (1997). The nature of precautionary wealth. *Journal of Monetary Economics* 40(1), 41–71. 9
- Carroll, C. D. and A. A. Samwick (1998). How important is precautionary saving? *Review of Economics and Statistics* 80(3), 410–419. 9

- Chetty, R. (2006). A general formula for the optimal level of social insurance. *Journal of Public Economics* 90(10), 1879–1901. 7, 41, 44
- Chetty, R. (2008). Moral hazard vs. liquidity and optimal unemployment insurance. *Journal of Political Economy* 116(2), 173–234. 1, 4, 1, 7, 7.1
- Chiappori, P. and B. Salanié (2000). Testing for asymmetric information in insurance markets. *Journal of Political Economy*, 56–78. 20
- Cullen, J. B. and J. Gruber (2000). Does unemployment insurance crowd out spousal labor supply? *Journal of Labor Economics* 18(3), 546–572. 41
- Dynan, K. E. (1993). How prudent are consumers? *Journal of Political Economy*, 1104–1113. 9
- Einav, L., A. Finkelstein, and J. Levin (2010, Sep). Beyond testing: Empirical models of insurance markets. *Annual Review of Economics* 2(1), 311–336. 1, 20
- Engen, E. M. and J. Gruber (2001). Unemployment insurance and precautionary saving. *Journal of Monetary Economics* 47(3), 545–579. 9, 41
- Feldstein, M. (1978). The effect of unemployment insurance on temporary layoff unemployment. *The American Economic Review*, 834–846. 41
- Finkelstein, A. and K. McGarry (2006). Multiple dimensions of private information: Evidence from the long-term care insurance market. *American Economic Review* 96(4), 938–958. 3
- Finkelstein, A. and J. Poterba (2002, Jan). Selection effects in the United Kingdom individual annuities market. *The Economic Journal*. 20
- Finkelstein, A. and J. Poterba (2004, Jan). Adverse selection in insurance markets: Policyholder evidence from the UK annuity market. *Journal of Political Economy*. 20
- Gan, L., M. Hurd, and D. McFadden (2005, Aug). Individual subjective survival curves. *Analyses in the Economics of Aging*, ed. D. Wise. 3, 4.2
- Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review* 87(1), 192–205. 1, 4, 4.3, 38, 5, 7.2

- Gruber, J. and J. B. Cullen (1996). Spousal labor supply as insurance: Does unemployment insurance crowd out the added worker effect? *NBER Working Paper*. 11, 7.1
- Guariglia, A. and B.-Y. Kim (2004). Earnings uncertainty, precautionary saving, and moonlighting in russia. *Journal of Population Economics* 17(2), 289–310. 9
- Guiso, L., T. Jappelli, and D. Terlizzese (1992). Earnings uncertainty and precautionary saving. *Journal of Monetary Economics* 30(2), 307–337. 9
- He, D. (2009, Jan). The life insurance market: Asymmetric information revisited. *Journal of Public Economics*. 3
- Hendren, N. (2013a). The policy elasticity. *NBER Working Paper*. 41, 43
- Hendren, N. (2013b). Private information and insurance rejections. *Econometrica* 81(5), 1713–1762. 1, 5, 6, 2.2, 18, 2.2, 3, 4.1, 1, 4.1, 4.1, 26, 4.1, 4.2, 4.2, 4.2, A.2
- Hendren, N. (2014). The inequality deflator: Interpersonal comparisons without a social welfare function. *NBER Working Paper*. 46
- Hubbard, R. G., J. Skinner, and S. P. Zeldes (1994). Precautionary saving and social insurance. *NBER Working Paper*. 9
- Kaplow, L. (2008). *The Theory of Taxation and Public Economics*. Princeton University Press. 46, 6.3
- Landais, C. (2015). Assessing the welfare effects of unemployment benefits using the regression kink design. *American Economic Journal: Economic Policy (Forthcoming)*. 4, 1
- Landais, C., P. Michailat, and E. Saez (2010). Optimal unemployment insurance over the business cycle. 7
- Lundberg, S. (1985). The added worker effect. *Journal of Labor Economics*, 11–37. 11
- Lusardi, A. (1997). Precautionary saving and subjective earnings variance. *Economics Letters* 57(3), 319–326. 9
- Lusardi, A. (1998). On the importance of the precautionary saving motive. *American Economic Review*, 449–453. 9

- Maloney, T. (1991). Unobserved variables and the elusive added worker effect. *Economica*, 173–187. 11
- Nekoei, A. and A. Weber (2015). Does extending unemployment benefits improve job quality? *CEPR Discussion Paper No. DP10568*. 41, 49
- Schmeider, J., T. von Wachter, and S. Bender (2013). The causal effect of unemployment duration on wages: Evidence from unemployment insurance extensions. *NBER Working Paper No. 19772*. 41, 49
- Shavell, S. (1979). On moral hazard and insurance. *The Quarterly Journal of Economics*, 541–562. 1, 2.2
- Shimer, R. and I. Werning (2007). Reservation wages and unemployment insurance. *The Quarterly Journal of Economics*, 1145–1185. 7
- Shimer, R. and I. Werning (2008). Liquidity and insurance for the unemployed. *The American Economic Review* 98(5), 1922–1942. 7
- Stephens, M. (2001). The long-run consumption effects of earnings shocks. *Review of Economics and Statistics* 83(1), 28–36. 1, 9
- Stephens, M. (2004). Job loss expectations, realizations, and household consumption behavior. *Review of Economics and statistics* 86(1), 253–269. 1, 10
- Topel, R. H. (1983). On layoffs and unemployment insurance. *The American Economic Review*, 541–559. 41

ONLINE APPENDIX: Not For Publication

A No Trade Condition

A.1 Multi-Dimensional Heterogeneity

This section considers the case in which there does not exist a one-to-one mapping between θ and $p(\theta)$ so that there is potentially heterogeneous willingness to pay for additional UI for different types θ with the same $p(\theta)$. I assume for simplicity that the distribution of $p(\theta)$ has full support on $[0, 1]$ and the distribution of $\frac{u'(c_e(\theta))}{v'(c_u(\theta))}$ has full support on $[0, \infty)$ (this is not essential, but significantly shortens the proof). I show that there exists a mapping, $f(p) : A \rightarrow \Theta$, where $A \subset [0, 1]$ such that the no trade condition reduces to testing

$$\frac{u'(c_u(f(p)))}{u'(c_e(f(p)))} \leq T(p) \quad \forall p$$

To see this, fix a particular policy, $\frac{d\tau}{db}$, and consider the set of θ that are willing to pay for this policy:

$$E \left[p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right]$$

Without loss of generality, there exists a function $\tilde{p} \left(\frac{d\tau}{db} \right)$ such that

$$E \left[p(\theta) \mid \theta \in \bar{\Theta} \left(\frac{d\tau}{db} \right) \right] = E \left[p(\theta) \mid p(\theta) \geq \tilde{p} \left(\frac{d\tau}{db} \right) \right]$$

so that the average probability of the types selecting $\frac{d\tau}{db}$ is equal to the average cost of types above $\tilde{p} \left(\frac{d\tau}{db} \right)$. Note that \tilde{p} can be constructed to be strictly increasing in $\frac{d\tau}{db}$ so that p^{-1} exists.

I construct $f(p)$ as follows. Define A to be the range of \tilde{p} when taking values of $\frac{d\tau}{db}$ ranging from 0 to ∞ . Without loss of generality, each value of $\frac{d\tau}{db}$ generates a different value of $p = \tilde{p}^{-1}(p)$. I assign $f(p)$ to each of these types as the value of θ such that

$$\frac{p}{1-p} \frac{u'(c_e(f(p)))}{v'(c_u(f(p)))} = \tilde{p}^{-1}(p)$$

which amounts to testing the no trade condition.

Intuitively, it is sufficient to check the no trade condition for the set of equivalent classes of types with the same willingness to pay for $\frac{d\tau}{db}$. Within this class, there exists a type that we can use to check the simple uni-dimensional no trade condition.

A.2 Robustness to Menus

Here, I illustrate how to nest the model into the setting of Hendren (2013b), then apply the no trade condition to rule out menus. Loosely, the present model is effectively the same model as in Hendren (2013b) aside from the introduction of a moral hazard problem and endogenous marginal utilities, u' and v' . I assume here that there is no heterogeneity in the marginal utilities across types and leave future work to study the problem of menus when individuals are making a range of additional choices. With this simplification, the only distinction relative to Hendren (2013b) is the introduction of the moral hazard problem in choosing p . Below, I show that introducing a moral hazard problem can't make trade any easier than in a world where $p(\theta)$ is exogenous and not affected by the insurer's contracts; hence the no trade condition results from Hendren (2013b) can be applied to rule out menus.

I abstract from individual heterogeneity in the utilities, $u(c)$ and $v(c)$, and assume for simplicity that consumption in the state of employment and unemployment is given exogenously and common across all types. I also assume individuals only choose p (i.e. there is no $a(\theta)$ choice). Introducing such behavioral responses and heterogeneity in utilities would likely be straightforward, but introduce a range of technical assumptions that would need to be included to rule out non-marginal insurance deviations. I leave a detailed treatment of this no trade condition under menus for future work.

I consider the maximization program of a monopolist insurer offering a menu of insurance contracts. Whether there exists any implementable allocations other than the endowment corresponds to whether there exists any allocations other than the endowment which maximize the profit, π , subject to the incentive and participation constraints.

The insurer can offer a menu of contracts, $\{\nu(\theta), \Delta(\theta)\}_{\theta \in \Gamma}$ where $\nu(\theta)$ specifies a total utility provided to type θ and $\Delta(\theta)$ denotes the difference in utilities if the agent becomes unemployed. Note that $\nu(\theta)$ implicitly contains the disutility of effort.

For exposition of the proof, I switch focus from the probability of unemployment, \hat{p} , to \hat{q} , which we define to be the probability of employment,

$$\hat{q}(\Delta; \theta) = 1 - \hat{p}(\Delta; \theta)$$

so that the agent's effort cost is $\Psi(\hat{q}(\Delta; \theta)\theta)$. Note that a type θ that accepts a contract containing Δ will choose a probability of employment $\hat{q}(\Delta; \theta)$ consistent with the first order

condition $\Psi'(\hat{q}(\Delta; \theta); \theta) = \Delta$.

Now, let $\pi(\Delta, \nu; \theta)$ denote the profits obtained from providing type θ with contract terms ν and Δ , given by

$$\pi(\Delta, \nu; \theta) = \hat{q}(\Delta; \theta)(c_e^e - C_e(v - \Psi(\Delta; \theta))) + (1 - \hat{q}(\Delta; \theta))(c_u^e - C_u(\nu - \Delta - \Psi(\Delta; \theta)))$$

where $C_u = u^{-1}$ and $C_e = v^{-1}$. Note that the profit function takes into account how the agents' choice of p varies with Δ .

Throughout, I maintain Assumption A.3: that π is concave in (ν, Δ) . Below in Section A.4, I discuss primitives for such concavity.

I prove the sufficiency of the no trade condition for ruling out trade by mapping it into the setting of Hendren (2013b) and applying his proof. To do so, define $\tilde{\pi}(\nu, \Delta; \theta)$ to be the profits incurred by the firm in the alternative world in which individuals choose p as if they faced their endowment (i.e. face no moral hazard problem). Now, in this alternative world, individuals still obtain total utility ν by construction, but must be compensated for their lost utility from effort because they can't re-optimize. But, note this compensation is second-order by the envelope theorem. Therefore, the marginal profitability for sufficiently small insurance contracts is given by

$$\pi(\nu, \Delta; \theta) \leq \tilde{\pi}(\nu, \Delta; \theta)$$

Now, define the aggregate profits to an insurer that offers menu $\{\nu(\theta), \Delta(\theta)\}_\theta$ by

$$\Pi(\nu(\theta), \Delta(\theta)) = \int \pi(\nu(\theta), \Delta(\theta); \theta) d\mu(\theta)$$

and in the world in which p is not affected by Π ,

$$\tilde{\Pi}(\nu(\theta), \Delta(\theta)) = \int \tilde{\pi}(\nu(\theta), \Delta(\theta); \theta) d\mu(\theta)$$

So, for small variations in ν and Δ , we have that

$$\Pi(\nu(\theta), \Delta(\theta)) \leq \tilde{\Pi}(\nu(\theta), \Delta(\theta))$$

because insurance causes an increase in p . Now, Hendren (2013b) shows that the no trade condition implies that $\tilde{\Pi} \leq 0$ for all menus, $\{\nu(\theta), \Delta(\theta)\}$. Therefore, the no trade condition also implies $\Pi \leq 0$ for local variations in the menu $\{\nu(\theta), \Delta(\theta)\}$ around the endowment. Combining with the concavity assumption, this also rules out larger deviations.

Conversely, if the no trade condition does not hold, note that the behavioral response is continuous in Δ , so that sufficiently small values of insurance allow for a profitable insurance contract to be traded.

A.3 Concavity Assumptions

Heretofore, I have placed no restrictions on either the nature of the distribution of types, θ , or the structure of the effort function, Ψ . This allows for considerable generality in characterizing when insurance markets can exist with moral hazard and adverse selection. But, the presence of moral hazard in this multi-dimensional screening problem induces the potential for non-convexities in the constraint set. Such non convexities could potentially limit the ability of local variational analysis to characterize the set of implementable allocations. Fortunately, a simple condition ensures that a local variational analysis provides a global characterization of the existence of profitable deviations from the endowment. Intuitively, the needed condition to ensure sufficiency of a local analysis is that the marginal profitability of insurance declines in the amount of insurance provided.

To express this condition, let Δ denote the difference in utilities between being employed and unemployed, so that lower values of Δ correspond to greater amounts of insurance. Define $\hat{p}(\Delta; \theta)$ to be the induced probability of unemployment for type θ , which solves

$$\Psi'(1 - \hat{p}(\Delta; \theta); \theta) = \Delta$$

It is straightforward to show that \hat{p} is decreasing in the size of the incentives to work, Δ . Now, define the cost functions,

$$\begin{aligned} C_u(x) &= u^{-1}(x) \\ C_v(x) &= v^{-1}(x) \end{aligned}$$

$C_u(x)$ measures the amount of consumption required to provide x units of utility when unemployed; similarly, $C_v(x)$ measures the amount of consumption required to provide x units of utility when employed.

Now, let $\pi(\Delta, \mu; \theta)$ denote the profit obtained from type θ if she is provided with total utility μ and difference in utilities Δ ,

$$\pi(\Delta, \mu; \theta) = (1 - \hat{p}(\Delta; \theta))(c_e^e - C_v(\mu - \Psi(1 - \hat{p}(\Delta; \theta)))) + \hat{p}(\Delta; \theta)(c_u^e - C_u(\mu - \Delta - \Psi(1 - \hat{p}(\Delta; \theta))))$$

To guarantee the validity of our variational analysis for characterizing when the endowment is the only implementable allocation, it will be sufficient to require that $\pi(\Delta, \mu; \theta)$ is concave in (Δ, μ) .

Assumption. $\pi(\Delta, \mu; \theta)$ is concave in (Δ, μ) for each θ

This assumption requires the marginal profitability of insurance to decline in the amount of insurance provided. If the agents choice of p is given exogenously (i.e. does not vary with Δ), then concavity of the utility functions, u and v , imply concavity of $\pi(\Delta, \mu; \theta)$. Assumption A.3 ensures that the ability of agents to choose p does not induce regions in which the marginal profitability of insurance actually increases in the amount of insurance.

A.4 Sufficient Conditions for Concavity

Assumption A.3 maintains that π is globally concave in (μ, Δ) . Here, we derive sufficient conditions on the primitives of the model that guarantee this concavity. In particular, we show that if $\Psi'''(q; \theta) > 0$ and $\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$ then π is globally concave in (μ, Δ) .

For simplicity, we consider a fixed θ and drop reference to it. Profits are given by

$$\pi(\Delta, \mu) = \hat{q}(\Delta)(c_e^e - C_e(\mu - \Psi(\hat{q}(\Delta)))) + (1 - \hat{q}(\Delta))(c_u^e - C_u(\mu - \Delta - \Psi(\hat{q}(\Delta))))$$

Our goal is to show the Hessian of π is negative semi-definite. We proceed in three steps. First, we derive conditions which guarantee $\frac{\partial^2 \pi}{\partial \Delta^2} < 0$. Second, we show that, in general, we have $\frac{\partial^2 \pi}{\partial \mu^2} < 0$. Finally, we show the conditions provided to guarantee $\frac{\partial^2 \pi}{\partial \Delta^2} < 0$ also imply the determinant of the Hessian is positive, so that both eigenvalues of the Hessian must be negative and thus the matrix is negative semi-definite.

A.4.1 Conditions that imply $\frac{\partial^2 \pi}{\partial \Delta^2} < 0$

Taking the first derivative with respect to Δ , we have

$$\begin{aligned} \frac{\partial \pi}{\partial \Delta} &= \frac{\partial \hat{q}}{\partial \Delta} (c_e^e - c_u^e + C_u(\mu - \Delta - \Psi(\hat{q}(\Delta)))) \\ &\quad - (1 - \hat{q}(\Delta)) C'_u(\mu - \Delta - \Psi(\hat{q}(\Delta))) - \hat{q}(\Delta) C'_e(\mu - \Psi(\hat{q}(\Delta))) \end{aligned}$$

Taking another derivative with respect to Δ , applying the identity $\Delta = \Psi'(\hat{p}(\Delta))$, and collecting terms yields

$$\begin{aligned} \frac{\partial^2 \pi}{\partial \Delta^2} = & - \left[(1 - \hat{q}(\Delta)) (1 + \Delta)^2 C''_u(\mu - \Delta - \Psi(\hat{q}(\Delta))) + \hat{q}(\Delta) (\Delta \hat{q}'(\Delta))^2 C''(\mu - \Psi(\hat{q}(\Delta))) \right] \\ & + \frac{\partial \hat{q}}{\partial \Delta} \left[(1 - \hat{q}(\Delta)) C'(\mu - \Delta - \Psi(\hat{q}(\Delta))) + \hat{q}(\Delta) C'(u - \Psi(\hat{q}(\Delta))) - (2 + 2\Delta \hat{q}'(\Delta)) C'(\mu - \Delta - \Psi(\hat{q}(\Delta))) \right] \\ & + \frac{\partial^2 \hat{q}}{\partial \Delta^2} \left[c_e^e - c_u^e + C(\mu - \Delta - \Psi(\hat{q}(\Delta))) - C(\mu - \Psi(\hat{q}(\Delta))) + (1 - \hat{q}(\Delta)) \Delta C'(\mu - \Delta - \Psi(\hat{q}(\Delta))) + \right. \end{aligned}$$

We consider these three terms in turn. The first term is always negative because $C'' > 0$. The second term, multiplying $\frac{\partial \hat{q}}{\partial \Delta}$, can be shown to be positive if

$$(1 + \hat{q}(\Delta)) C'(\mu - \Delta - \Psi(\hat{q}(\Delta))) \geq \hat{q}(\Delta) C'(\mu - \Delta)$$

which is necessarily true whenever

$$\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$$

This inequality holds as long as people are willing to pay less than a 100% markup for a small amount of insurance, evaluated at their endowment.

Finally, the third term is positive as long as $\Psi''' > 0$. To see this, one can easily verify that the term multiplying $\frac{\partial^2 \hat{q}}{\partial \Delta^2}$ is necessarily positive. Also, note that $\frac{\partial^2 \hat{q}}{\partial \Delta^2} = \frac{-\Psi'''}{(\Psi'')^2}$. Therefore, if we assume that $\Psi''' > 0$, the entire last term will necessarily be negative. In sum, it is sufficient to assume $\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$ and $\Psi''' > 0$ to guarantee that $\frac{\partial^2 \pi}{\partial \Delta^2} < 0$.

A.4.2 Conditions that imply $\frac{\partial^2 \pi}{\partial \mu^2} < 0$

Fortunately, profits are easily seen to be concave in μ . We have

$$\frac{\partial \pi}{\partial \mu} = - (1 - \hat{q}(\Delta)) C'(\mu - \Delta - \Psi(\hat{q}(\Delta))) - \hat{q}(\Delta) C'(\mu - \Psi(\hat{q}(\Delta)))$$

so that

$$\frac{\partial^2 \pi}{\partial \mu^2} = - (1 - \hat{q}(\Delta)) C''(\mu - \Delta - \Psi(\hat{q}(\Delta))) - \hat{q}(\Delta) C''(\mu - \Psi(\hat{q}(\Delta)))$$

which is negative because $C'' > 0$.

A.4.3 Conditions to imply $\frac{\partial^2 \pi}{\partial \Delta^2} \frac{\partial^2 \pi}{\partial \mu^2} - \left(\frac{\partial^2 \pi}{\partial \Delta \partial \mu} \right) > 0$

Finally, we need to ensure that the determinant of the Hessian is positive. To do so, first note that

$$\frac{\partial^2 \pi}{\partial \mu \partial \Delta} = (1 - \hat{q}(\Delta)) C''(\mu - \Delta - \Psi(\hat{q}(\Delta))) (1 + \Delta \hat{q}'(\Delta)) + \hat{q}(\Delta) C''(\mu - \Psi(\hat{q}(\Delta))) \Delta \hat{q}'(\Delta)$$

Also, we note that under the assumptions $\Psi''' > 0$ and $\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$, we have the inequality

$$\frac{\partial^2 \pi}{\partial \Delta^2} < - \left[(1 - \hat{q}(\Delta)) (1 + \Delta)^2 C_u''(\mu - \Delta - \Psi(\hat{q}(\Delta))) + \hat{q}(\Delta) (\Delta \hat{q}'(\Delta))^2 C''(\mu - \Psi(\hat{q}(\Delta))) \right]$$

Therefore, we can ignore the longer terms in the expression for $\frac{\partial^2 \pi}{\partial \Delta^2}$ above. We multiply the RHS of the above equation with the value of $\frac{\partial^2 \pi}{\partial \mu^2}$ and subtract $\left(\frac{\partial^2 \pi}{\partial \Delta \partial \mu} \right)^2$. Fortunately, many of the terms cancel out, leaving the inequality

$$\begin{aligned} \frac{\partial^2 \pi}{\partial \Delta^2} \frac{\partial^2 \pi}{\partial \mu^2} - \left(\frac{\partial^2 \pi}{\partial \Delta \partial \mu} \right)^2 &\geq (1 - \hat{q}(\Delta)) \hat{q}(\Delta) (1 + \Delta \hat{q}'(\Delta))^2 C_u''(\mu - \Delta - \Psi(\hat{q}(\Delta))) C''(\mu - \Psi(\hat{q}(\Delta))) \\ &\quad + \hat{q}(\Delta) (1 - \hat{q}(\Delta)) (\Delta \hat{q}'(\Delta))^2 C''(\mu - \Psi(\hat{q}(\Delta))) C''(\mu - \Delta - \Psi(\hat{q}(\Delta))) \\ &\quad - 2(1 - \hat{q}(\Delta)) \hat{q}(\Delta) (1 + \Delta \hat{q}'(\Delta)) \Delta \hat{q}'(\Delta) C_u''(\mu - \Delta - \Psi(\hat{q}(\Delta))) C''(\mu - \Psi(\hat{q}(\Delta))) \end{aligned}$$

which reduces to the inequality

$$\frac{\partial^2 \pi}{\partial \Delta^2} \frac{\partial^2 \pi}{\partial \mu^2} - \left(\frac{\partial^2 \pi}{\partial \Delta \partial \mu} \right)^2 \geq \hat{q}(\Delta) (1 - \hat{q}(\Delta)) C_u''(\mu - \Delta - \Psi(\hat{q}(\Delta))) C''(\mu - \Psi(\hat{q}(\Delta))) K(\mu, \Delta)$$

where

$$\begin{aligned} K(\mu, \Delta) &= (1 + \Delta \hat{q}'(\Delta))^2 + (\Delta \hat{q}'(\Delta))^2 - 2\Delta \hat{q}'(\Delta) - 2(\Delta \hat{q}'(\Delta))^2 \\ &= 1 \end{aligned}$$

So, since $C'' > 0$, we have that the determinant must be positive. In particular, we have

$$\frac{\partial^2 \pi}{\partial \Delta^2} \frac{\partial^2 \pi}{\partial \mu^2} - \left(\frac{\partial^2 \pi}{\partial \Delta \partial \mu} \right)^2 \geq \hat{q}(\Delta) (1 - \hat{q}(\Delta)) C_u''(\mu - \Delta - \Psi(\hat{q}(\Delta))) C''(\mu - \Psi(\hat{q}(\Delta)))$$

A.4.4 Summary

As long as $\Psi''' > 0$ and $\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$, the profit function is guaranteed to be concave. As noted in the text, generally one finds empirically that $\frac{u'(c_u^e)}{v'(c_e^e)} \leq 2$. Therefore, the unsubstantiated assumption for the model is that the convexity of the effort function increases in p , $\Psi''' > 0$. An alternative statement of this assumption is that $\frac{\partial^2 \hat{q}}{\partial \Delta^2} < 0$, so that the marginal impact of work incentives on the employment probability is declining in the size of the work incentives.

Future work can derive the necessary conditions when the willingness to pay for additional UI varies conditional on p and when individuals can make additional actions, $a(\theta)$, in response to unemployment. I suspect the proofs can be extended to such cases, but identifying the necessary conditions for global concavity would be an interesting direction for future work.

B Proof of Proposition 1

I prove the proposition in two steps. First, I show that $\text{cov}\left(P, \frac{m(P)}{P}\right) \leq 0$. Then, I use the Lemma to

Lemma 1. For any P , it must be the case that $\text{cov}\left(P, \frac{m(P)}{P}\right) \leq 0$.

Proof: note that

$$m(P) = E[P - p | P \geq p]$$

so that

$$\text{cov}\left(P, \frac{m(P)}{P}\right) = E[m(P)] - E[P] E\left[\frac{m(P)}{P}\right]$$

So, we wish to show that

$$\frac{E[m(P)]}{E[P]} < E\left[\frac{m(P)}{P}\right]$$

Note that:

$$E\left[\frac{m(P)}{P}\right] = E\left[\frac{\frac{1}{1-F(P)} \int (\tilde{p} - P) f(\tilde{p}) d\tilde{p}}{P}\right] = E\left[\frac{E[\tilde{p} | \tilde{p} \geq P]}{P}\right] - 1 = E_P E_{\tilde{p}}\left[\frac{\tilde{p}}{P} | \tilde{p} \geq P\right] - 1$$

And:

$$\frac{E[m(P)]}{E[P]} = \frac{E_P E_{\tilde{p}}[\tilde{p} | \tilde{p} \geq P]}{E[P]} - 1$$

So, we wish to test whether

$$E_P E_{\tilde{p}}\left[\frac{\tilde{p}}{E[P]} | \tilde{p} \geq P\right] <? E_P E_{\tilde{p}}\left[\frac{\tilde{p}}{P} | \tilde{p} \geq P\right]$$

or

$$E_P E_{\tilde{p}}\left[\frac{\tilde{p}}{P} - \frac{\tilde{p}}{E[P]} | \tilde{p} \geq P\right] >? 0$$

or

$$E_P E_{\tilde{p} \geq P}\left[\tilde{p} \left(\frac{1}{P} - \frac{1}{E[P]}\right) | \tilde{p} \geq P\right] >? 0$$

Note that once we've conditioned on $\tilde{p} \geq P$, we can replace \tilde{p} with P and maintain an inequality

$$\begin{aligned} E_P E_{\tilde{p} \geq P}\left[\tilde{p} \left(\frac{1}{P} - \frac{1}{E[P]}\right) | \tilde{p} \geq P\right] &\geq E_P E_{\tilde{p} \geq P}\left[P \left(\frac{1}{P} - \frac{1}{E[P]}\right) | \tilde{p} \geq P\right] \\ &\geq E_P E_{\tilde{p} \geq P}\left[1 - \frac{P}{E[P]} | \tilde{p} \geq P\right] \\ &\geq E_P\left[1 - \frac{P}{E[P]}\right] \\ &\geq 0 \end{aligned}$$

Which implies $\text{cov}\left(\frac{m(P)}{P}, P\right) < 0$.

Proof of Proposition.

Note that since $E[P|P \geq p] \geq p$,

$$\begin{aligned} E[T(P)] &= E_p \left[\frac{E[P|P \geq p]}{p} \frac{1-p}{1-E[P|P \geq p]} \right] \\ &\geq E_p \left[1 + \frac{m(p)}{p} \right] \end{aligned}$$

So, it suffices to show that $E\left[\frac{m(P)}{P}\right] \geq \frac{E[m(P)]}{E[P]}$. Clearly

$$E[m(P)] = E\left[\frac{m(P)}{P}\right] E[P] + \text{cov}\left(P, \frac{m(P)}{P}\right)$$

so that

$$E\left[\frac{m(P)}{P}\right] = \frac{E[m(P)] - \text{cov}\left(P, \frac{m(P)}{P}\right)}{E[P]}$$

by Lemma 1, $\text{cov}\left(P, \frac{m(P)}{P}\right) \leq 0$. So,

$$E\left[\frac{m(P)}{P}\right] \geq \frac{E[m(P)]}{E[P]} = \frac{E[m(P)]}{\Pr\{U\}}$$

so that

$$E[T(P)] \geq E\left[1 + \frac{m(P)}{P}\right] \geq 1 + \frac{E[m(P)]}{\Pr\{U\}}$$

which is the desired result.

C Welfare Metrics

C.1 Ex-post consumption derivation

Euler equation when $\beta R = 1$:

$$v'(c_{pre}(p)) = pu'(c_u(p)) + (1-p)v'(c_e(p))$$

so that

$$pu'(c_u(p)) = v'(c_{pre}(p)) - (1-p)v'(c_e(p))$$

Now,

$$\frac{E\left[\frac{p}{p}u'(c_u(p))\right]}{E\left[\frac{1-p}{1-p}v'(c_e(p))\right]} = \frac{\frac{1}{p}E[v'(c_{pre}(p)) - (1-p)v'(c_e(p))]}{E\left[\frac{1-p}{1-p}v'(c_e(p))\right]}$$

Now, we can take a Taylor expansion around $\bar{c}_e = E[c_e]$ yielding

$$\begin{aligned}
\frac{\frac{1}{\bar{p}} E[v'(c_{pre}(p)) - (1-p)v'(c_e(p))]}{E\left[\frac{1-p}{1-\bar{p}}v'(c_e(p))\right]} &\approx \frac{\frac{1}{\bar{p}} E[v' + v''(c_{pre}(p) - \bar{c}_e) - (1-p)(v' + v''(c_e(p) - \bar{c}_e))]}{v' + v'' \text{cov}\left(\frac{1-p}{1-\bar{p}}, c_e(p)\right)} \\
&= \frac{\frac{1}{\bar{p}} [\bar{p}v' + v''[(\bar{c}_{pre} - \bar{c}_e) - v'' E[(1-p)(c_e(p) - \bar{c}_e)]]]}{v' + v'' \text{cov}\left(\frac{1-p}{1-\bar{p}}, c_e(p)\right)} \\
&= \frac{1 + \frac{1}{\bar{p}} \left[(-\sigma) \left[\left(\frac{\bar{c}_{pre} - \bar{c}_e}{\bar{c}_e}\right) + E\left[(1-p) \left(\frac{c_e(p) - \bar{c}_e}{\bar{c}_e}\right)\right]\right]\right]}{1 + (-\sigma) \text{cov}\left(\frac{1-p}{1-\bar{p}}, \frac{c_e(p)}{\bar{c}_e}\right)} \\
&= \frac{1 + \frac{1}{\bar{p}} \left[\sigma \left[\frac{\Delta c^{good}}{c} + \text{cov}\left[(1-p), \left(\frac{c_e(p) - \bar{c}_e}{\bar{c}_e}\right)\right]\right]\right]}{1 + (-\sigma) \text{cov}\left(\frac{1-p}{1-\bar{p}}, \frac{c_e(p)}{\bar{c}_e}\right)}
\end{aligned}$$

which is roughly

$$1 + \frac{1}{\bar{p}} \sigma \frac{\Delta c^{good}}{c} + \sigma \text{cov}\left(\frac{-p}{\bar{p}}, \frac{c_e(p)}{\bar{c}_e}\right)$$

where

$$\frac{\Delta c^{good}}{c} = \frac{\bar{c}_e - \bar{c}_{pre}}{\bar{c}_e}$$

This uses the variation in c in the employed state. Under the assumption that not losing your job increases earnings, it implies that

$$W \geq \sigma \text{cov}\left(\frac{-p}{\bar{p}}, \frac{c_e(p)}{\bar{c}_e}\right) \approx \sigma \frac{\text{var}(Z) \text{cov}(-Z, \log(c_e))}{\Pr\{L\} \text{var}(Z)}$$

which can be obtained with regression.

C.2 Ex-ante labor supply derivation

We also observe labor supply responses by households in response to these shocks. If we assume a spouse labor supply decision, $l \in \{0, 1\}$, is contained in the set of other actions, a . Suppose this earns income, y . Then, we can use the spousal labor supply response, combined with known estimates of the spousal labor response to labor earnings to back out the implied value of social insurance. Let

$$\Psi(1-p, a, \theta) = \tilde{\Psi}(1-p, \tilde{a}, \theta) + 1\{l=1\} \eta(\theta)$$

where $\eta(\theta)$ is the disutility of labor for type θ , distributed F_η in the population.

Let $k(y, l, p)$ denote the value to a type p of choosing l to obtain income y when they face an unemployment probability of p . The labor supply decision is

$$k(y, 1, p) - k(0, 0, p) \geq \eta(\theta)$$

so that types will choose to work if and only if it increases their utility. This defines a threshold rule whereby individuals choose to work if and only if $\eta(\theta) \leq \bar{\eta}(y, p)$ and the labor force participation rate is given by $\Phi(y, p) = F(\bar{\eta}(y, p))$.

Now, note that

$$\frac{d\Phi}{dp} = f(\bar{\eta}) \frac{\partial \bar{\eta}}{\partial p} = f(\bar{\eta}) \left[\frac{\partial k(y, 1, p)}{\partial p} - \frac{\partial k(0, 0, p)}{\partial p} \right]$$

and making an approximation that the impact of the income y does not discretely change the instantaneous marginal utilities (i.e. because it will be smoothed out over the lifetime or because the income is small), we have

$$\frac{d\Phi}{dp} \approx f(\bar{\eta}) \frac{\partial^2 k}{\partial p^2} y$$

Finally, note that $\frac{\partial k}{\partial y} = v'(c_{pre}(p))$ is the marginal utility of income. So,

$$\frac{d\Phi}{dp} \approx f(\bar{\eta}) \frac{d}{dp} [v'(c_{pre}(p))] y$$

and integrating across all the types p , we have

$$E_p \left[\frac{d\Phi}{dp} \right] \approx E_p \left[f(\bar{\eta}) \frac{d}{dp} v'(c_{pre}(p)) y \right]$$

To compare this response to a wage elasticity, consider the response to a \$1 increase in wages

$$\frac{d\Phi}{dy} = f(\bar{\eta}) \frac{\partial k}{\partial y}$$

so,

$$E_p \left[\frac{d\Phi}{dp} \right] \approx E_p \left[\frac{d\Phi}{dy} y \frac{\frac{d}{dp} v'(c_{pre}(p))}{v'(c_{pre}(p))} \right]$$

Now, let $\epsilon^{semi} = \frac{d\Phi}{d \log(y)}$ denote the semi-elasticity of spousal labor force participation. We therefore have

$$\frac{E_p \left[\frac{d\Phi}{dp} \right]}{\epsilon^{semi}} \approx E_p \left[\frac{\frac{d}{dp} v'(c_{pre}(p))}{v'(c_{pre}(p))} \right]$$

so that the ratio of the labor supply response to p divided by the semi-elasticity of labor supply with respect to wages reveals the average elasticity of the marginal utility function. Assuming this elasticity is roughly constant and noting that a Taylor expansion suggests that for any function $f(x)$, we have $\frac{f(1)-f(0)}{f(0)} \approx \frac{d}{dx} \log(f)$, we have

$$\frac{E_p \left[\frac{d\Phi}{dp} \right]}{\epsilon^{semi}} \approx \frac{v'(1) - v'(0)}{v'(0)}$$

Now, how do we estimate $\frac{d\Phi}{dp}$? We see $\Phi(Z)$. regressing l on Z will generate an attenuated coefficient. To first order, we can inflate this by the ratio of the variance of Z to the variance of P , or

$$\frac{v'(1) - v'(0)}{v'(0)} \approx \beta \frac{1}{\epsilon^{semi}} \frac{var(Z)}{var(P)}$$

TABLE I
Sample Summary Statistics

Variable	Panel 1: Baseline Sample		Panel 2: Health Sample		Panel 3: Consumption Sample	
	mean	std dev	mean	std dev	mean	std dev
<i>Selected Observables (subset of X)</i>						
Age	56.1	5.1	56.1	5.2	56.8	4.8
Male	0.40	0.49	0.41	0.49	0.36	0.48
Wage	36,057	143,883	37,523	154,993	40,523	59,402
Job Tenure (Years)	12.7	10.8	12.7	10.9	12.9	10.7
<i>Unemployment Outcome (U)</i>	0.031	0.173	0.032	0.175	0.036	0.187
<i>Subjective Probability Elicitation (Z)</i>	15.7	24.8	15.7	24.6	15.7	23.8
<i>Sample Size</i>						
Number of Observations		26,640		22,831		2,798
Number of Households		3,467		3,180		862

Notes: This table presents summary statistics for the samples used in the paper. Panel 1 presents the baseline sample used in Part I of the analysis. Panel 2 presents the statistics for the subset of the baseline sample that have non-zero health variables for the extended controls used in Part I. Panel 3 presents the summary statistics for the ~10% sub-sample of the baseline sample who took part in the consumption survey component of the HRS. The rows present selected summary statistics, including the age of respondents, gender, wage, and job tenure. The unemployment outcome is defined using the subsequent survey wave to construct an indicator for the individual losing his/her job involuntarily in the subsequent 12 months after the baseline survey.

TABLE II
Lower Bound Estimates

Specification:	Alternative Controls			Sub-Samples					
	Baseline (1)	Demo (2)	Health (3)	Age ≤ 55 (4)	Age > 55 (5)	Below Median Wage (6)	Above Median Wage (7)	Tenure > 5 yrs (8)	Tenure ≤ 5 yrs (9)
$E[T_Z(P_Z)-1]$	0.7682	# 0.8033	0.7198	0.6983	0.8150	0.6513	0.9515	1.0996	0.4759
s.e.	(0.053)	# (0.051)	(0.052)	(0.081)	(0.066)	(0.058)	(0.094)	(0.093)	(0.057)
$E[m_Z(P_Z)]$	0.0236	# 0.0247	0.0228	0.0208	0.0256	0.0244	0.0228	0.0192	0.0270
s.e.	(0.002)	# (0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)	(0.002)	(0.003)
p-value	0.0000	# 0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
$\Pr\{U=1\}$	0.0307	# 0.0307	0.0317	0.0297	0.0314	0.0375	0.0239	0.0175	0.0568
Controls									
Demographics	X	X	X	X	X	X	X	X	X
Job Characteristics	X		X	X	X	X	X	X	X
Health			X						
Num of Obs.	26640	26640	22831	11134	15506	13320	13320	17681	8959
Num of HHs	3467	3467	3180	2255	3231	2916	2259	2939	2447

Notes: Table presents estimates of the nonparametric lower bounds on $E[T(P)]$ and the average mean residual risk function, $E[m(P)]$. Column (1) reports the baseline specification. Columns (2) uses only demographic controls; Column (3) uses demographic, job characteristics, and health characteristics. Columns (4)-(9) report results for the baseline specification on various subsamples including below and above age 55 (Columns 4 and 5), above and below-median wage earners (Columns 6 and 7) and above and below 5 years of job tenure. The first row presents estimates of the lower bounds of $E[T(P)]$, which is computed as $1+E[m_Z(P_Z)]/\Pr\{U=1\}$. The value of $E[m_Z(P_Z)]$ is reported in the second row. This is computed using the distribution of predicted values (illustrated in Figure II, Panel A). I construct the average predicted value above a given threshold within an age-by-gender aggregation window; Appendix Table I reports the robustness to alternative aggregation windows. The third row reports the p-value from the test that the coefficients in the probit specification for $\Pr\{U|X,Z\}$ are all equal to zero, clustering the standard errors at the household level. All standard errors for $E[T_Z(P_Z)]$ and $E[m_Z(P_Z)]$ are constructed using 500 bootstrap resamples at the household level.

TABLE III
Minimum Pooled Price Ratio

Specification	Alternative Controls			Sub-Samples					
	Baseline (1)	Demo (2)	Health (3)	Age <= 55 (4)	Age > 55 (5)	Below Median Wage (6)	Above Median Wage (7)	Tenure > 5 yrs (8)	Tenure <= 5 yrs (9)
Inf T(p) - 1	3.360	5.301	3.228	3.325	3.442	4.217	3.223	4.736	3.739
s.e.	(0.203)	(0.655)	(0.268)	(0.306)	(0.279)	(0.417)	(0.268)	(0.392)	(0.336)
Controls									
Demographics	X	X	X	X	X	X	X	X	X
Job Characteristics	X		X	X	X	X	X	X	X
Health Characteristics			X						
Num of Obs.	26,640	26,640	22,831	11,134	15,506	13,320	13,320	17,850	8,790
Num of HHs	3,467	3,467	3,180	2,255	3,231	2,916	2,259	2,952	2,437

Notes: This table presents estimates of the minimum pooled price ratio, $\text{inf T}(p)$. Column (1) reports the baseline specification. Columns (2) uses only demographic controls; Column (3) uses demographic, job characteristics, and health characteristics. Columns (4)-(9) report results for the baseline specification on various subsamples including below and above age 55 (Columns 4 and 5), above and below-median wage earners (Columns 6 and 7) and above and below 5 years of job tenure. The table reports the minimum pooled price ratio across the 3 point masses included in the distribution, excluding the highest value of the point mass (which is mechanically 1). Appendix Table III provides the estimated distribution values. All parameter estimates are constructed using maximum likelihood. Because of the non-convexity of the optimization program, I assess the robustness to 1000 initial starting values. All standard errors are constructed using bootstrap re-sampling using 1000 re-samples at the household level.

TABLE IV
Consumption Relationship to Potential Job Loss

<i>Specification:</i>	Baseline	HH Cons	Sample Z > 0	Sample U = 0	Non-Durable Consumption	No Controls	No 1 {Z=0} Control
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Panel 1: Estimation of $d[\log(c)]/dZ$</i>							
Elicitation (Z)	-0.160**	-0.110*	-0.171**	-0.162**	-0.162**	-0.345***	-0.0401
s.e.	(0.0781)	(0.0596)	(0.0777)	(0.0783)	(0.0789)	(0.0798)	(0.0659)
Elicitation * Unemp (Z*U)	-0.137	-0.421**	-0.0771		-0.257	-0.0000475	-0.460**
s.e.	(0.268)	(0.207)	(0.268)		(0.303)	(0.296)	(0.218)
Elicitation of 0 (1 {Z=0})	-0.0893***	-0.0587**		-0.0904***	-0.120***	-0.160***	
s.e.	(0.0334)	(0.0279)		(0.0334)	(0.0356)	(0.0365)	
Elicitation of 0 * Unemp (1 {Z=0}*U)	0.338	0.161			0.307	0.191	
s.e.	(0.222)	(0.180)			(0.220)	(0.239)	
Unemp (U)	-0.0845	0.0862	-0.120		-0.0936	-0.181	0.118
s.e.	(0.165)	(0.128)	(0.164)		(0.164)	(0.187)	(0.120)

<i>Panel 2: Welfare Calculation</i>							
Scaling Factor (Var(Z X) / Pr{L})	1.53	1.53	1.20	1.40	1.53	1.55	1.53
bootstrap s.e.	(0.15)	(0.15)	(0.14)	(0.15)	(0.14)	(0.16)	(0.15)
Cov(p/E[p],log(c _e))	0.25**	0.17*	0.21**	0.23**	0.25**	0.54***	0.06
bootstrap s.e.	(0.12)	(0.09)	(0.10)	(0.11)	(0.12)	(0.13)	(0.10)
Implied WTP (CRRA = 2)	0.49**	0.34*	0.41**	0.45**	0.5**	1.07***	0.12
bootstrap s.e.	(0.23)	(0.19)	(0.19)	(0.23)	(0.23)	(0.27)	(0.20)

Mean Dep Var	9.86	10.58	9.89	9.87	9.18	9.86	9.86
Num of Obs.	2,798	2,798	1,503	2,696	2,798	2,798	2,798
Num of HHs	862	862	579	843	862	862	862

Notes: This table presents estimates from a regression of log consumption expenditure on subjective elicitation of becoming unemployed and indicators of the event of actually becoming unemployed in the subsequent 12 months. Consumption expenditure is measured 12 months after the subjective probability elicitation, and asks about consumption expenditure covering the previous 12 months. Columns (1) and (3)-(7) use log household consumption per capita as the dependent variable, taking the household consumption expenditure and dividing it by the total number of household members before taking the log. Column (2) uses log household consumption. Column (1) reports the baseline results for a specification that includes the elicitation, Z, an indicator for Z=0 to capture the nonlinearity in Figure IV, an indicator for subsequent unemployment, U, an interaction of the elicitation with the indicator for unemployment, and an interaction of an indicator for Z=0 with the indicator for future unemployment, U. Column (2) replicates Column (1) with household consumption as the dependent variable. Column (3) restricts the sample to those with positive elicitation. Column (4) restricts the sample to those who do not become unemployed in the subsequent 12 months (U=0). Column (5) replicates the specification in Column (1) using non-durable consumption per capita instead of total consumption per capita. Column (6) drops all control variables for age, gender, log wage, year, and region. Column (7) considers the specification in Column (1) but drops the indicators for focal point responses at Z=0.

Panel 1 reports the baseline coefficients from the regressions. Panel 2 reports the implied welfare implications. The first row reports the scaling factor, $\text{var}(Z|X)/\text{Pr}\{L\}$, and its bootstrapped standard error. The second row multiplies the regression coefficient on the elicitation by this scaling factor to construct the estimate for $\text{Cov}(p, E[p], \log(c_e))$. The third row presents the implied willingness to pay under an assumption of the coefficient of relative risk aversion over consumption in the employed state of 2, which is simply $2 * \text{Cov}(p, E[p], \log(c_e))$. Standard errors are constructed via 500 bootstrap samples, resampling at the household level.

TABLE V
Spousal Labor Supply Relationship to Potential Job Loss

Specification:	Entry								
	Baseline with Z=0 Intercept	No 1 {Z=0} Control	Sample without Future Job Loss	Full Time Work	2yr Lagged Entry ("Placebo")	Household Fixed Effects	Individual Fixed Effects	Exit	Spouse Unemployment
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>Panel 1: Estimation of dL/dZ</i>									
Elicitation (Z)	0.0273** (0.0112)	0.0282*** (0.00868)	0.0270** (0.0116)	0.0286** (0.0128)	0.00792 (0.0102)	0.0267* (0.0146)	0.0312 (0.0230)	0.0170 (0.0116)	0.0250*** (0.00964)
s.e.									
Elicitation of 0 (1 {Z=0})	-0.000605 (0.00465)		-0.000492 (0.00472)	0.000635 (0.00554)	0.00237 (0.00486)	0.000246 (0.00598)	0.00191 (0.0102)		
<i>Panel 2: Welfare Calculation</i>									
Total Var / Signal Var (var(Z X)/var(P X))	11.00 (1.41)	11.00 (1.40)	11.00 (1.37)	11.00 (1.38)	18.17 (3.90)	11.00 (1.32)	11.00 (1.32)		
bootstrap s.e.									
$W_{\text{welfare}} (\epsilon^{\text{semi}} = 0.5)$	0.6** (0.26)	0.62*** (0.21)	0.59** (0.26)	0.63** (0.30)	0.29 (0.41)	0.59** (0.29)	0.69* (0.39)		
bootstrap s.e.									
Mean Dep Var	0.04	0.04	0.04	0.05	0.04	0.04	0.04	0.09	0.03
Num of Obs.	11049	11049	10726	11049	11049	11049	11049	11049	9079
Num of HHs	2214	2214	2194	2214	2214	2214	2214	2214	1359

Notes: This table presents the coefficients from a regression of spousal labor entry on the subjective elicitation. I restrict the sample to respondents who are married in both the current and previous wave. I define spousal entry as an indicator for the event that both (a) the spouse was not working for pay in the previous wave (2 years prior) and (b) the spouse is currently working for pay. For Columns (1)-(7) I include observations for which the spouse was working for pay in the previous wave (these observations are coded as zero). Column (1) presents a linear regression of an indicator for spousal labor entry on the elicitation, Z, and controls for age, age squared, gender, log wage, year, and census division (10 regions), and an indicator for Z=0 to deal with potential non-linearities resulting from focal point responses. Column (2) drops the indicator for Z=0. Column (3) restricts to the subsample that does not lose their job in the subsequent 12 months. Column (4) defines spousal labor force entry using only full time employment. I define an indicator for the event that both (a) the spouse was not employed full time in the previous wave and (b) is currently working full time. Column (5) uses the lagged value of Z from the previous wave (2 years prior) as a "placebo" test. Note this is not an exact placebo test to the extent to which the information is correlated across time. Column (6) adds household fixed effects to the specification in Column (1). Column (7) adds individual fixed effects to the specification in Column (1). Column (8) replaces the dependent variable with an indicator for exit of the spouse from the labor market. I define exit as an indicator for being in the labor force last wave (2 years prior) and out of the labor force this wave. Column (9) replaces the dependent variable with an indicator for spouse unemployment in the subsequent 12 months and restricts the sample to spouses currently in the labor market.

Panel 2 presents the welfare implications of each model. I scale the regression coefficient in Panel 1 by the total variance of Z relative to the signal variance (var(P)). I estimate the variance of Z given X by regressing Z on the control variables and squaring the RMSE. I estimate the variance of P given X as follows. I regress the future unemployment indicator, U, on the controls and take the residuals. I regress Z on the controls and take those residuals. I then construct the covariance between these two residuals and rescale by $(n-1)/(n-df)$, where df is the number of degrees of freedom in the regression of U on the controls. This provides an estimate of $\text{Cov}(Z, U|X)$, which is an approximation to $\text{var}(P|X)$ that is exact under classical measurement error. The Implied WTP is constructed by taking the regression coefficient, multiplying by the total variance / signal variance, and dividing by the semi-elasticity of spousal labor supply, here assumed to be 0.5. For example, the 0.6 in Column (1) is obtained by $0.0273 * 11 / 0.5 = 0.60$. All standard errors in Panel 2 are constructed using 500 bootstrap repetitions, resampling at the household level.

TABLE VI
Impact of Future Job Loss on Consumption and 2-Sample IV Welfare Estimates

Specification:	Full Sample (1)	Employed in t-2 and t-1 (2)	Controls for Needs (3)	Individual Fixed Effects (4)	Over 40 Sample (5)	With Outliers (6)
<i>Panel 1: Impact of Unemployment on $\log(c_{t-2}) - \log(c_{t-1})$</i>						
Unemp	-0.0336*** (0.00570)	-0.0250*** (0.00942)	-0.0249** (0.00994)	-0.0231* (0.0130)	-0.0287* (0.0151)	-0.0231* (0.0121)
s.e.						
<i>Specification Details</i>						
Sample Employed in t-2 and t-1		X	X	X	X	X
Controls for change in log needs (t-2 vs t-1)			X	X	X	
Individual Fixed Effects				X		
Mean Dep Var	0.049	0.053	0.054	0.054	0.036	0.060
Num of Obs.	80984	70503	58987	58987	27264	72758
Num of HHs	11055	10042	8869	8869	4772	10156

<i>Panel 2: Split-Sample IV Welfare Calculation</i>						
Δ _{First Stage}	0.103 (0.016)	0.103 (0.016)	0.103 (0.016)	0.103 (0.016)	0.103 (0.016)	0.103 (0.016)
bootstrap s.e.						
$d[\log(c_{\text{pec}}(p))]/dp$ (2-sample 2SLS)	0.33*** (0.06)	0.24*** (0.09)	0.24*** (0.10)	0.22* (0.13)	0.28* (0.15)	0.22* (0.12)
s.e.						
$W^{\text{ex-ante}}$ ($\sigma = 2$)	0.65*** (0.11)	0.48*** (0.18)	0.48*** (0.19)	0.45* (0.25)	0.56* (0.29)	0.45* (0.23)
s.e.						

Notes: This Table presents the split-sample IV estimates of the impact of p on log consumption. The sample includes all household heads Panel 1 reports the coefficients from a regression of the change in log food consumption between years t-2 and t-1 on an indicator of unemployment. Column (1) controls for a cubic in age and year dummies. Column (2) restricts the sample to those who are not unemployed in either t-2 or t-1. Column (3) adds controls for the change in log expenditure needs ("need_std_p") between t-2 and t-1 and the change in total household size between t-2 and t-1. Column (4) adds individual fixed effects to the specification in Column (3). Column (5) restricts the sample to those 40 and over for the specification in Column (3). Following Gruber (1997), Columns (1)-(5) drop observations with more than a 3-fold change in consumption; Column (6) includes these outliers following the specification for those who are not unemployed in both t-1 and t-2. Panel 2 reports the impact of unemployment on the elicitation. The first row reports the difference in the coefficient from a regression of the elicitation, Z , on subsequent unemployment in the next year, U , and the coefficient from a regression of Z on an indicator for unemployment in the 12-24 months after the survey. Appendix Table IV provides the baseline regression results for this first stage calculation. The standard error is computed using bootstrap resampling at the household level (500 repetitions). The consumption drop equivalent reports divides the coefficient in Panel 1 by the coefficient on the regression of the elicitation on unemployment to arrive at the estimate of $d\log(c)/dp$. The Implied WTP multiplies this by a coefficient of relative risk aversion of 2. Standard errors for the consumption drop equivalent are approximated as the reduced form standard errors scaled by the coefficient on U in the first stage. All standard errors are clustered at the household level.

APPENDIX TABLE I
Alternative Lower Bound Specifications

Specification:	Baseline		Alternative Aggregation Windows			Alternative Subj. Prob Spec			Alternative Outcomes		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	
$E[m_Z(P_Z)-1]$	0.7687	0.6802	0.7716	0.7058	0.7150	0.7462	0.7681	0.5296	0.3675	0.5790	
s.e.	(0.058)	(0.051)	(0.05)	(0.048)	(0.048)	(0.051)	(0.054)	(0.033)	(0.04)	(0.086)	
$E[m_Z(P_Z)]$	0.0239	0.0209	0.0237	0.0217	0.0220	0.0229	0.0236	0.0314	0.0147	0.0067	
s.e.	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	
p-value	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	
$Pr\{U=1\}$	0.0310	0.0307	0.0307	0.0307	0.0307	0.0307	0.0307	0.0593	0.0401	0.0115	
Controls											
Demographics	X	X	X	X	X	X	X	X	X	X	
Job Characteristics	X	X	X	X	X	X	X	X	X	X	
Elicitation Specification											
Polynomial Degree	3	3	3	3	3	1	4	3	3	3	
Focal pt dummies (0, 50, 100)	X	X	X	X	X	X	X	X	X	X	
Aggregation Window	Age x Gender	Age x Gender	Constant	Age x Gender x Industry	Age x Gender x Occupation	Age x Gender	Age x Gender	Age x Gender	Age x Gender	Age x Gender	
Unemployment Outcome Window	12 months	12 months	12 months	12 months	12 months	12 months	12 months	24 months	6-24 months	6-12 months	
Error Specification	Probit	Linear	Probit	Probit	Probit	Probit	Probit	Probit	Probit	Probit	
Num of Obs.	25516	26640	26640	26640	26640	26640	26640	26640	26640	26640	
Num of HHs	3467	3467	3467	3467	3467	3467	3467	3467	3467	3467	

Notes: Table reports robustness of lower bound estimates in Table II to alternative specifications. Column (1) replicates the baseline specification in Table II (Column (1)). Column (2) constructs the predicted values, $Pr\{U|X,Z\}$, using a linear model instead of a probit specification. Columns (3)-(5) consider alternative aggregation windows for translating the distribution of predicted values into estimates of $E[m_Z(P_Z)]$. While Column (1) constructs $m_Z(P_Z)$ using the predicted values within age-by-gender groups, Column (3) aggregates the predicted values across the entire sample. Column (4) uses a finer partition, aggregating within age-by-gender-by-industry groups. Column (5) aggregates within age-by-gender-by-occupation groups. Columns (6)-(7) consider alternative specifications for the subjective probability elicitation. Column (6) uses only a linear specification in Z combined with focal point indicators at Z=0, Z=50, and Z=100, as opposed to the baseline specification that also includes a polynomial in Z. Column (7) adds a third and fourth order polynomial in Z to the baseline specification. Columns (8)-(10) consider alternative outcome definitions for U. Column (8) defines unemployment, U, as an indicator for involuntary job loss at any point in between survey waves (24 months). Column (9) defines unemployment as an indicator for job loss in between survey waves excluding the first six months after the survey (i.e. 6-24 months). Finally, Column (10) defines unemployment as an indicator for job loss in the 6-12 months after the survey wave.

APPENDIX TABLE II
Estimation of $F(p|X)$

Specification	Alternative Controls			Sub-Samples					
	Baseline	Demo	Health	Age <= 55	Age > 55	Below Median Wage	Above Median Wage	Tenure > 5 yrs	Tenure <= 5 yrs
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
1st mass									
Location	0.001	0.012	0.002	0.001	0.002	0.007	0.000	0.000	0.022
s.e.	(0.001)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.001)	(0.000)	(0.003)
Weight	0.446	0.713	0.449	0.437	0.461	0.530	0.452	0.422	0.612
s.e.	(0.024)	(0.071)	(0.054)	(0.035)	(0.030)	(0.032)	(0.034)	(0.036)	(0.034)
T(p)	63.839	6.301	39.032	101.038	36.986	12.413	262.088	6.9E+08	5.052
s.e.	6.1E+06	1.7E+00	1.8E+06	1.0E+07	1.1E+06	3.2E+00	7.6E+07	2.5E+08	6.0E-01
2nd mass									
Location	0.031	0.031	0.032	0.030	0.031	0.037	0.024	0.018	0.0575
s.e.	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Weight	0.471	0.202	0.470	0.483	0.456	0.365	0.486	0.508	0.2771
s.e.	(0.024)	(0.071)	(0.052)	(0.035)	(0.030)	(0.032)	(0.034)	(0.037)	(0.0341)
T(p)	4.360	8.492	4.228	4.325	4.442	5.217	4.223	5.736	4.7392
s.e.	0.203	4.194	4.576	0.306	0.279	0.417	2.181	3.008	0.5227
3rd Mass									
Location	0.641	0.639	0.642	0.639	0.643	0.626	0.649	0.641	0.6420
s.e.	(0.004)	(0.004)	(0.028)	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.0055)
Weight	0.082	0.086	0.081	0.081	0.083	0.105	0.061	0.070	0.1105
s.e.	(0.002)	(0.002)	(0.006)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.0040)
Controls									
Demographics	X	X	X	X	X	X	X	X	X
Job Characteristics	X		X	X	X	X	X	X	X
Health Characteristics			X						
Num of Obs.	26,640	26,640	22,831	11,134	15,506	13,320	13,320	17,850	8,790
Num of HHs	3,467	3,467	3,180	2,255	3,231	2,916	2,259	2,952	2,437

Notes: This table presents estimates of the distribution of private information about unemployment risk, P . Column (1) reports the baseline specification. Columns (2) uses only demographic controls; Column (3) uses demographic, job characteristics, and health characteristics. Columns (4)-(9) report results for the baseline specification on various subsamples including below and above age 55 (Columns 4 and 5), above and below-median wage earners (Columns 6 and 7) and above and below 5 years of job tenure. The $F(p)$ estimates report the location and mass given to each point mass, evaluated at the mean $q = \Pr\{U=1\} = 0.031$. For example, in the baseline specification, the results estimate a point mass at 0.001, 0.031, and 0.641 with weights 0.446, 0.471 and 0.082. The values of $T(p)$ represent the markup that individuals at this location in the distribution would have to be willing to pay to cover the pooled cost of worse risks. All parameter estimates are constructed using maximum likelihood. Because of the non-convexity of the optimization program, I assess the robustness to 1000 initial starting values. All standard errors are constructed using bootstrap re-sampling using 1000 re-samples at the household level.

APPENDIX TABLE III
Summary Statistics (PSID Sample)

	mean	std dev
<i>Variable</i>		
Age	39.794	10.27
Male	0.808	0.39
Unemployment	0.059	0.24
Year	1985	7.62
Log Consumption	8.199	0.65
Log Expenditure Needs	8.124	0.32
 <i>Consumption growth ($\log(c_{t-2})-\log(c_{t-1})$)</i>	 0.049	 0.360
 <i>Sample Size</i>		
Number of Observations	80,984	
Number of Households	11,055	

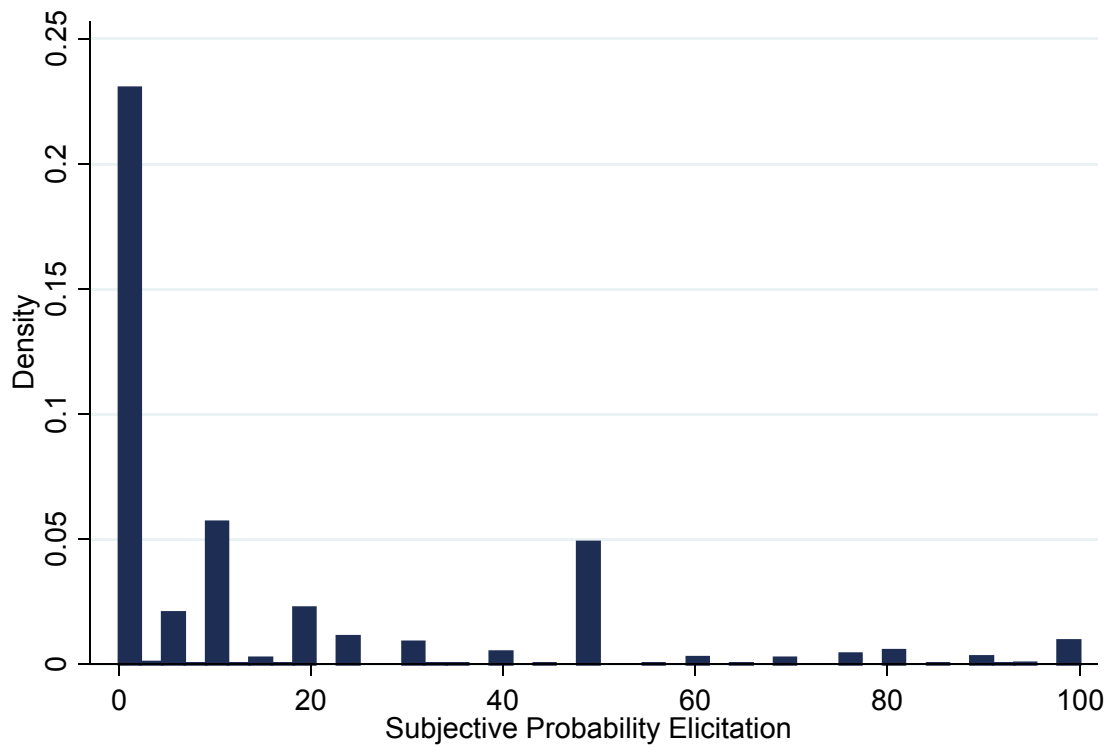
Notes: This table presents the summary statistics for the PSID sample used to estimate the impact of future unemployment on consumption growth in the year prior to unemployment. I use data from the PSID for years 1971-1997. Sample includes all household heads with non-missing variables.

APPENDIX TABLE IV
Information Realization Between t-2 and t-1 ("First Stage")

	Coeff. on Z
	(1)
<i>Dependent Variable:</i>	
Unemp (Next 12 months)	0.197***
s.e.	(0.0123)
Unemp (12-24 months)	0.0937***
s.e.	(0.0113)
Difference	0.1031***
bootstrap s.e.	(0.0159)
Num of Obs.	26,640


Note: This table presents estimates from two separate regressions of the elicitation on unemployment in the subsequent 12 months (U) and the elicitation, Z, on unemployment in the 12-24 months after the elicitation. The standard error for the difference is computed using 500 bootstrap repetitions resampling at the household level.

FIGURE I: Histogram of Subjective Probability Elicitations



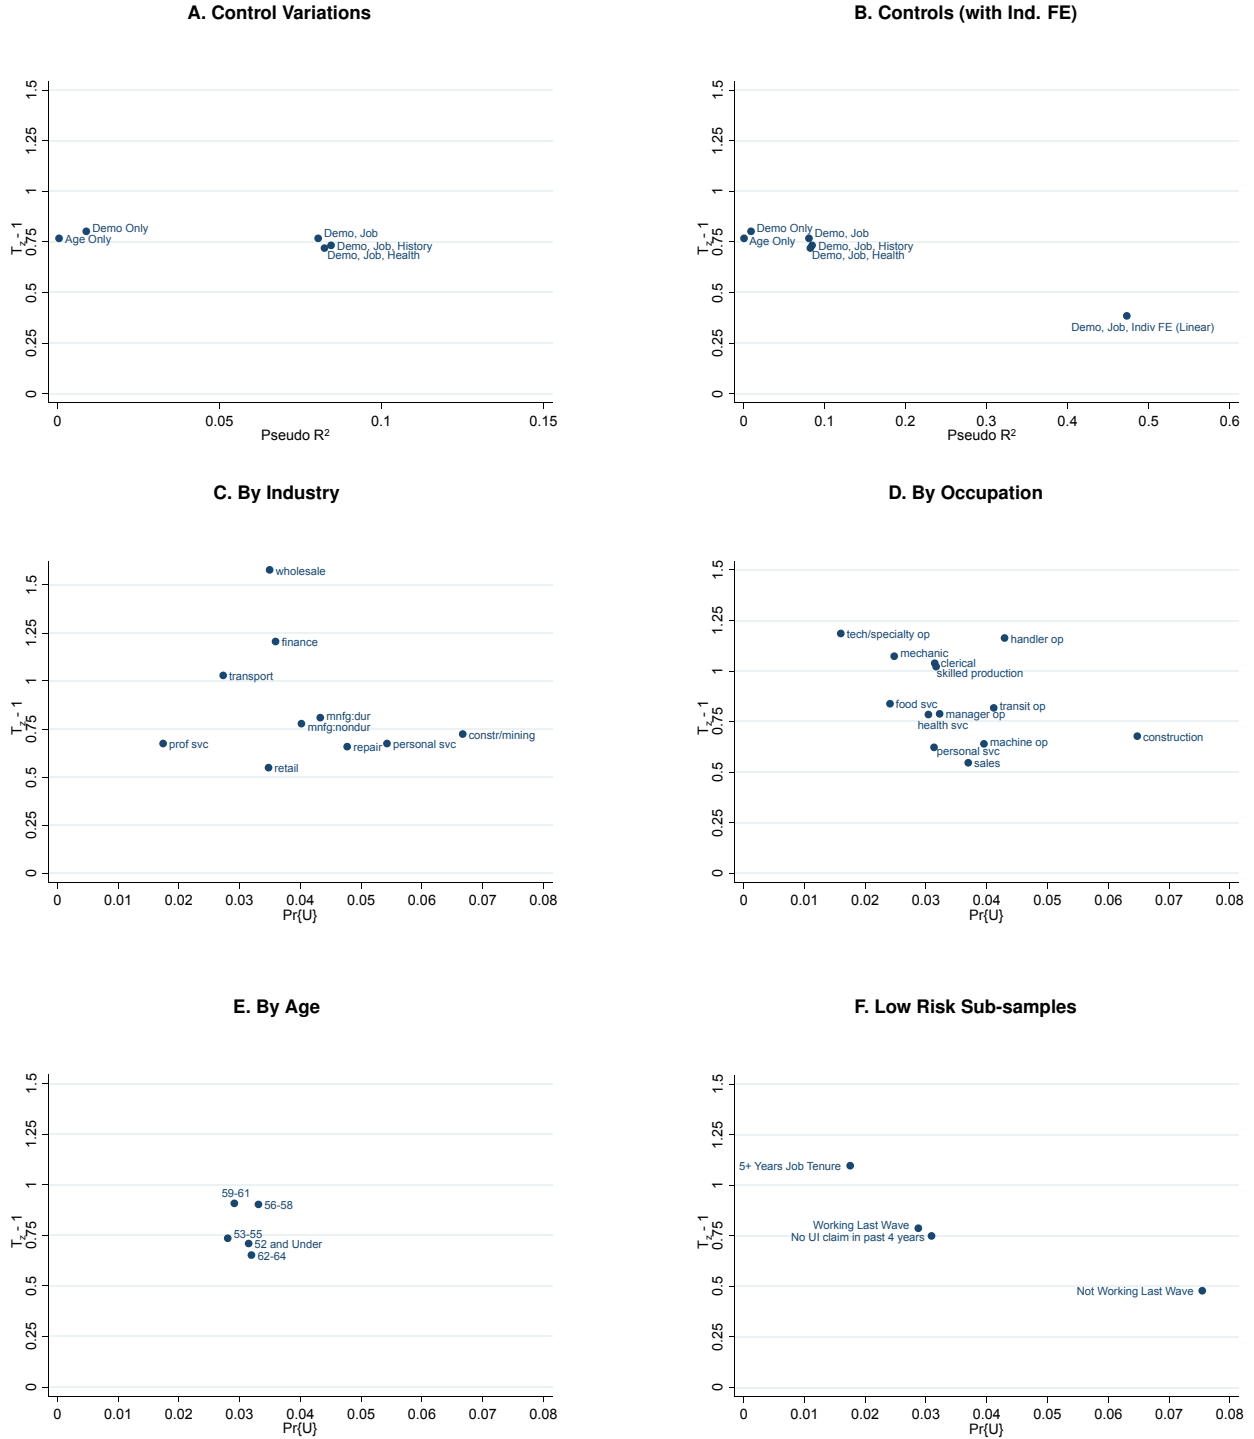
Notes: This figure presents a histogram of responses to the question “What is the percent chance (0-100) that you will lose your job in the next 12 months?”. The figure reports the histogram of responses for the baseline sample outlined in Panel 1 of Table I. As noted in previous literature, responses tend to concentrate on focal point values, especially $Z = 0$.

FIGURE II: Predictive Content of Subjective Probability Elicitations




Notes: These figures present the predictive content in the subjective probability elicitation. Panel A reports the mean unemployment rate in each elicitation category controlling for demographic and job characteristics. To construct this figure, I first regress the unemployment indicator on the demographic and job characteristics and take the residuals. I then construct the mean of these residuals in each of the elicitation categories and add back the mean unemployment rate. To obtain the 5 / 95% confidence intervals, I run a regression of unemployment on each of these categories with zero as the omitted category, clustering the standard errors by household. Panel B reports the kernel density of the distribution of predicted values from a regression of both observables and the elicitation on U , $\Pr\{U|X, Z\}$, minus the predicted values from a regression of U on observables, X , $\Pr\{U|X\}$. Under the Assumptions outlined in the text, the true distribution of P given X is a mean-preserving spread of this distribution of predicted values.

FIGURE III: Lower Bounds for $E[T(P)]$



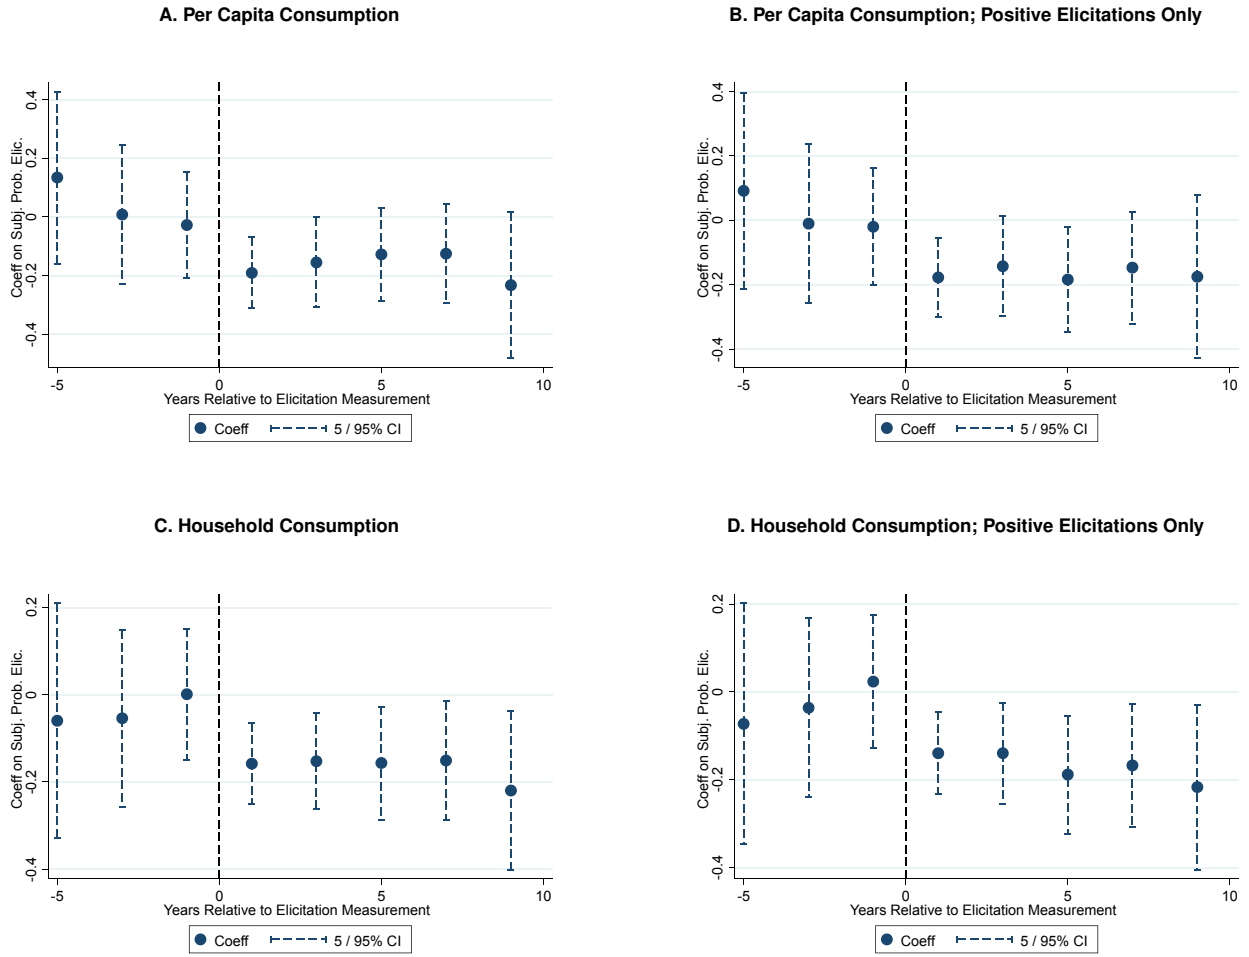
Notes: These figures present estimates of the lower bounds on the average pooled price ratio, $E[T_Z(P_Z)]$, using a range of sub-samples and controls. Panel A reports estimates of $E[T_Z(P_Z)]$ for a range of control variables. Panel B adds a specification with individual fixed effects to Panel A and relies on a linear specification as opposed to a probit (see Appendix Table I, Column (2) for the baseline estimation using the linear model). The horizontal axis presents the Pseudo- R^2 of the specification for $\Pr\{U|X, Z\}$. Panel C constructs separate estimates by industry classification. Panel D constructs estimates by occupation. Panel E constructs separate estimates for each wave of the survey. Panel F restricts the sample to varying sub-samples, analyzing the relationship between $E[T_Z(P_Z)]$ and restrictions to lower-risk subsamples. The horizontal axis in Panels C-F report the mean unemployment probability, $\Pr\{U\}$, for each sub-sample.

FIGURE IV: Relationship between Potential Job Loss and Consumption



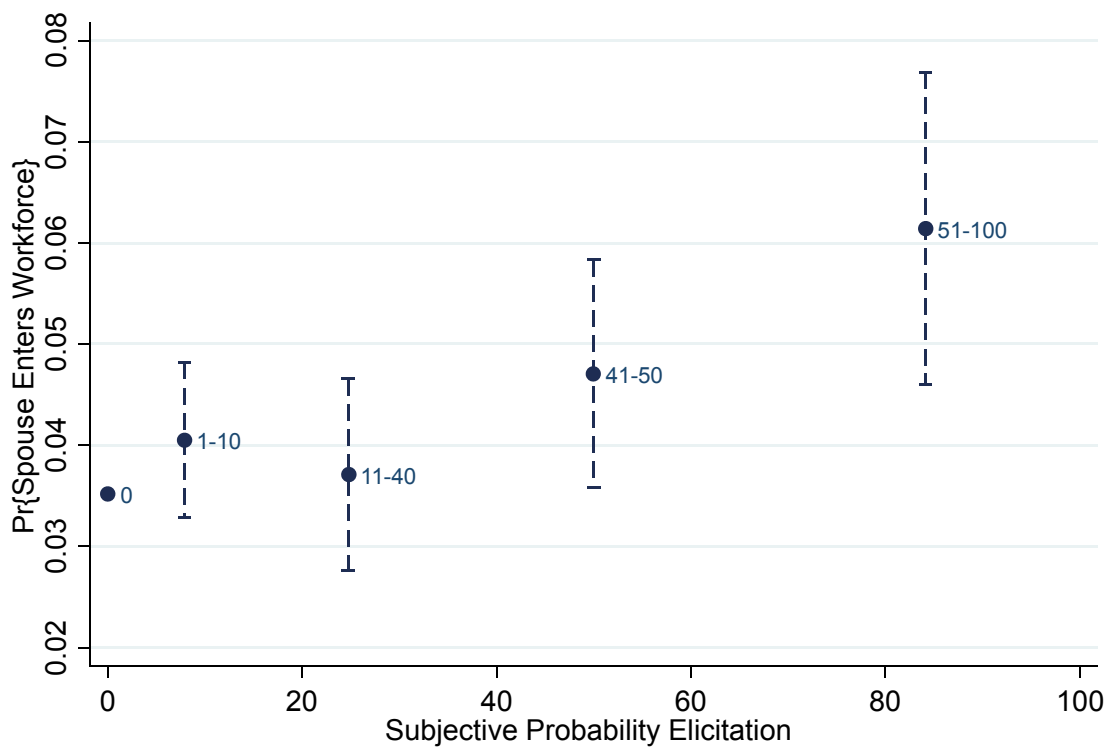
Notes: These figures present coefficients from a regression of log household consumption per capita (Panel A) and log total household consumption (Panel B) on category indicators for the subjective probability elicitation, Z , controlling for realized unemployment status, U , and several observable characteristics: age, age squared, gender, year dummies, census division, log wage, and an indicator for being married.

FIGURE V: Relationship between Potential Job Loss and Consumption Leads and Lags




Notes: These figures present coefficients from a regression of leads and lags of log per capita consumption (Panels A and B) and log household consumption (Panels C and D) on the subjective probability elicitation, controlling for an indicator for realized unemployment, an indicator for a subjective probability elicitation of $Z = 0$, and several observable characteristics: age, age squared, gender, year dummies, census division, log wage, and an indicator for being married. Panels A and C include all observations; Panels B and D restrict the sample to those with positive elicitation, $Z > 0$. The vertical dotted line corresponds to the time of the subjective probability elicitation. The horizontal axis corresponds to the time of the consumption measurement (which includes a 12 month look-back window).

FIGURE VI: Relationship between Potential Job Loss and Spousal Labor Supply




Notes: The figure present coefficients from a regression of an indicator for a spouse entering the labor force – defined as an indicator for not working in the previous wave and working in the current wave – on category indicators for the subjective probability elicitation, Z , controlling for realized unemployment status, U , and several observable characteristics: age, age squared, gender, year dummies, census division, log wage, and an indicator for being married.

FIGURE VII: Impact of Unemployment on Consumption Growth




Notes: These figures present coefficients from separate regressions of leads and lags of the log change in food expenditure on an indicator of unemployment, along with controls for year indicators and a cubic in age. Sample is restricted to household heads. Food expenditure is the sum of food in the home, food outside the home, and food stamps. Following Gruber (1997) and Chetty et al. (2005), I define food stamps by taking the monthly measure and multiplying by 12 for the years where the monthly food stamp measure is available. The horizontal axis presents the years of the lead/lag for the consumption expenditure growth measurement (i.e. 0 corresponds to consumption growth in the year of the unemployment measurement relative to the year prior to the unemployment measurement). Panel A presents the results for the full sample. Panel B restricts the sample to household heads who are not unemployed in $t - 1$ or $t - 2$.

ONLINE APPENDIX FIGURE I: Additional Lower Bounds on $E[T(P)]$



Notes: This figure presents additional estimates of the lower bound on the average pooled price ratio, $E[T_Z(P_Z)]$. Panel A reports separate estimates for each wave of the survey and Panel B reports estimates by census division. Panel C reports a set of estimates that use alternative definitions of U . This includes an indicator for involuntarily losing one's job for three time windows: in between surveys (0-24 months), in the 6-12 months after the survey, and 6-24 months after the survey. The 6-12 and 6-24 month specifications simulate lower bounds on $E[T_Z(P_Z)]$ in a hypothetical underwriting scenario whereby insurers would impose 6 month waiting periods. I also include specifications that interact these indicators with indicators that the individual had positive government UI claims, which effectively restricts to the subset of unemployment spells where the individual takes up government UI benefits.

ONLINE APPENDIX FIGURE II: Comparison to Other Non-Existing Insurance Markets



Notes: Hendren (2013) argues private information prevents people with pre-existing conditions from purchasing insurance in LTC, Life, and Disability insurance markets. This figure compares the estimates of $\inf T(p) - 1$ for the baseline specification in the unemployment context to the estimates in Hendren (2013) for the sample of individuals who are unable to purchase insurance due to a pre-existing condition. Figure reports the confidence interval and the 5 / 95% confidence interval for each estimate in each sample.